


Prometheus Operator:

"Out-of-the-Box" Monitoring As A Code

Щербаков Станислав Николаевич



HighLoad++
Весна 2021



About Me

Профессиональные навыки

- 25 лет уверенный пользователь ПК
- Более 10 лет работы по специальности
- ~5 лет из них DevOps Engineer

Хобби

- INFOSEC
- Reverse engineering
- Console hacking



t.me/xSTASiANx

About Plesk

- 400k инсталляций продукта
- 6% сайтов в интернете (NetCraft)
- Более 20 лет на рынке
- 150+ классных инженеров

plesk


Table of contents

- The way from Legacy monitoring systems to Modern
- Monitoring Problems
- Legacy Infrastructure Service Discovery
- Prometheus Operator as the Solution
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

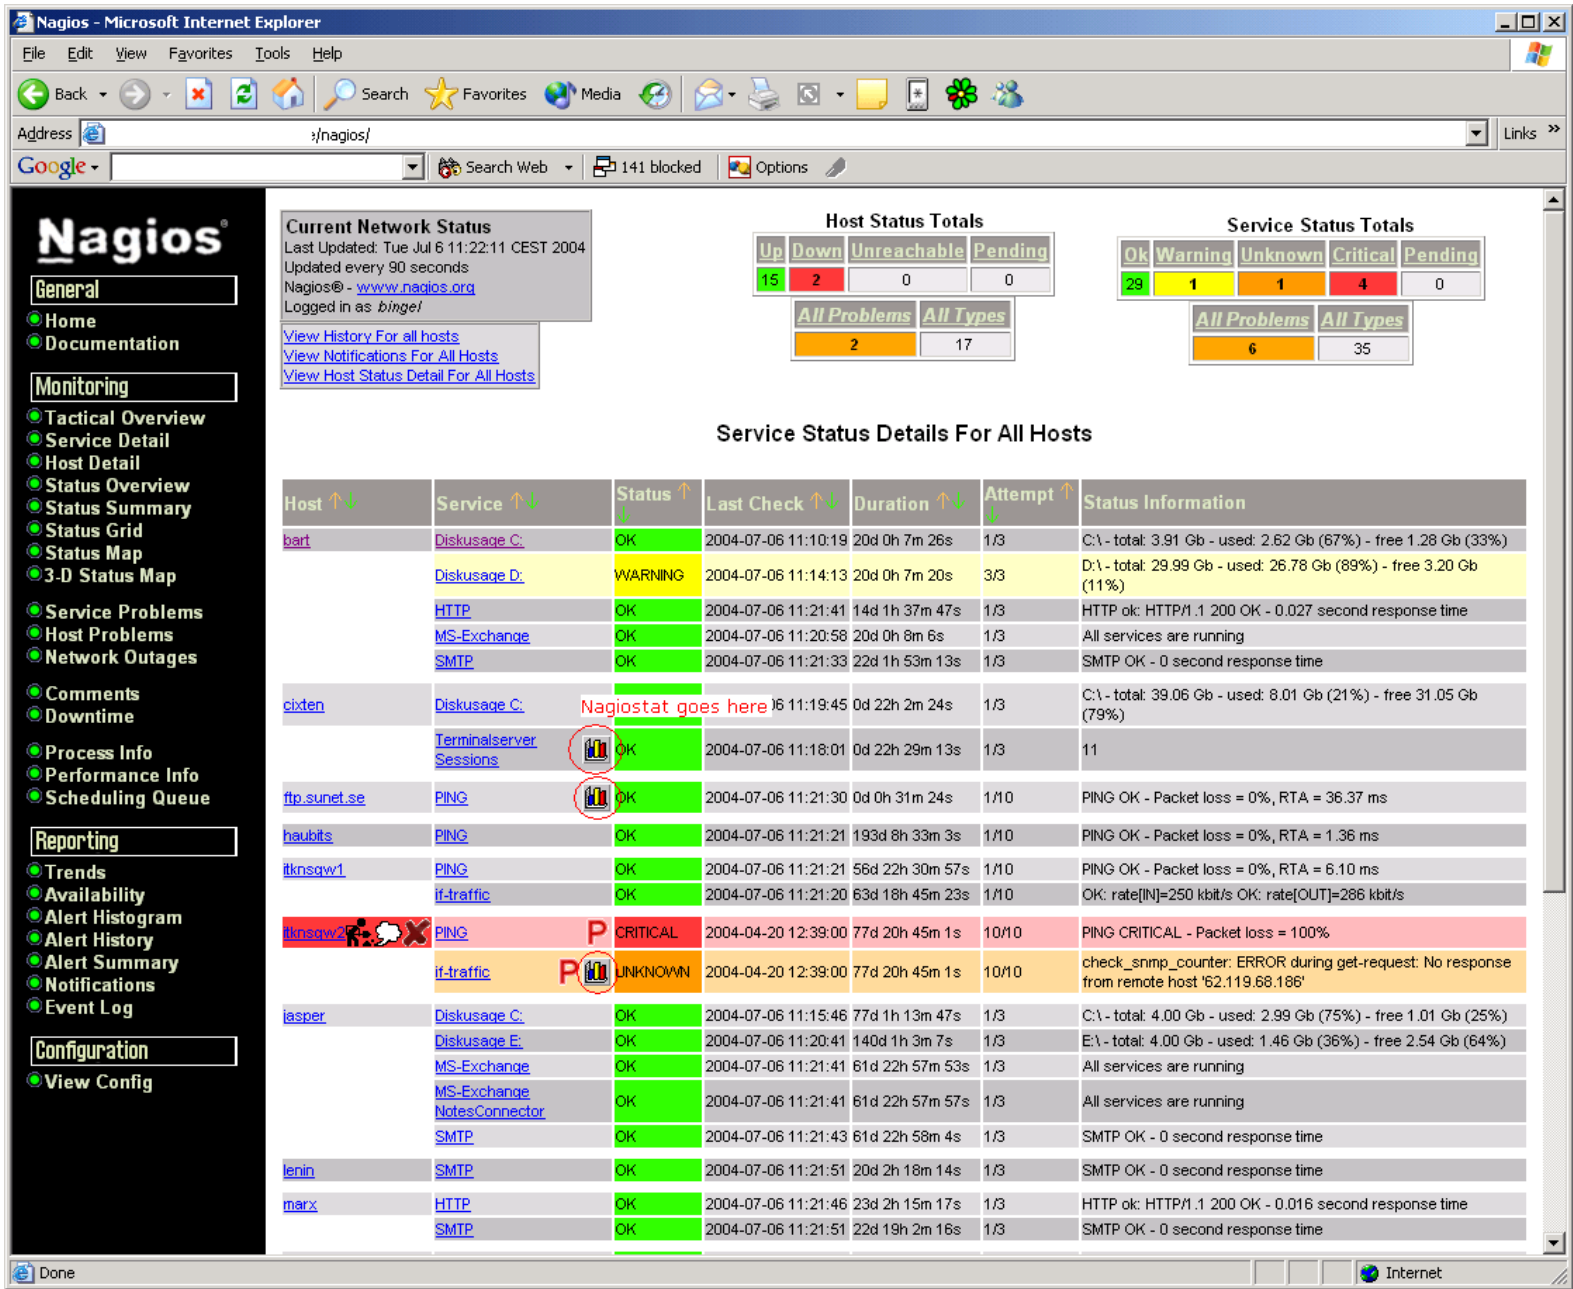
Now we are here

- The way from Legacy monitoring systems to Modern
- Monitoring Problems
- Legacy Infrastructure Service Discovery
- Prometheus Operator as the Solution
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

Monitoring systems




Legacy monitoring systems

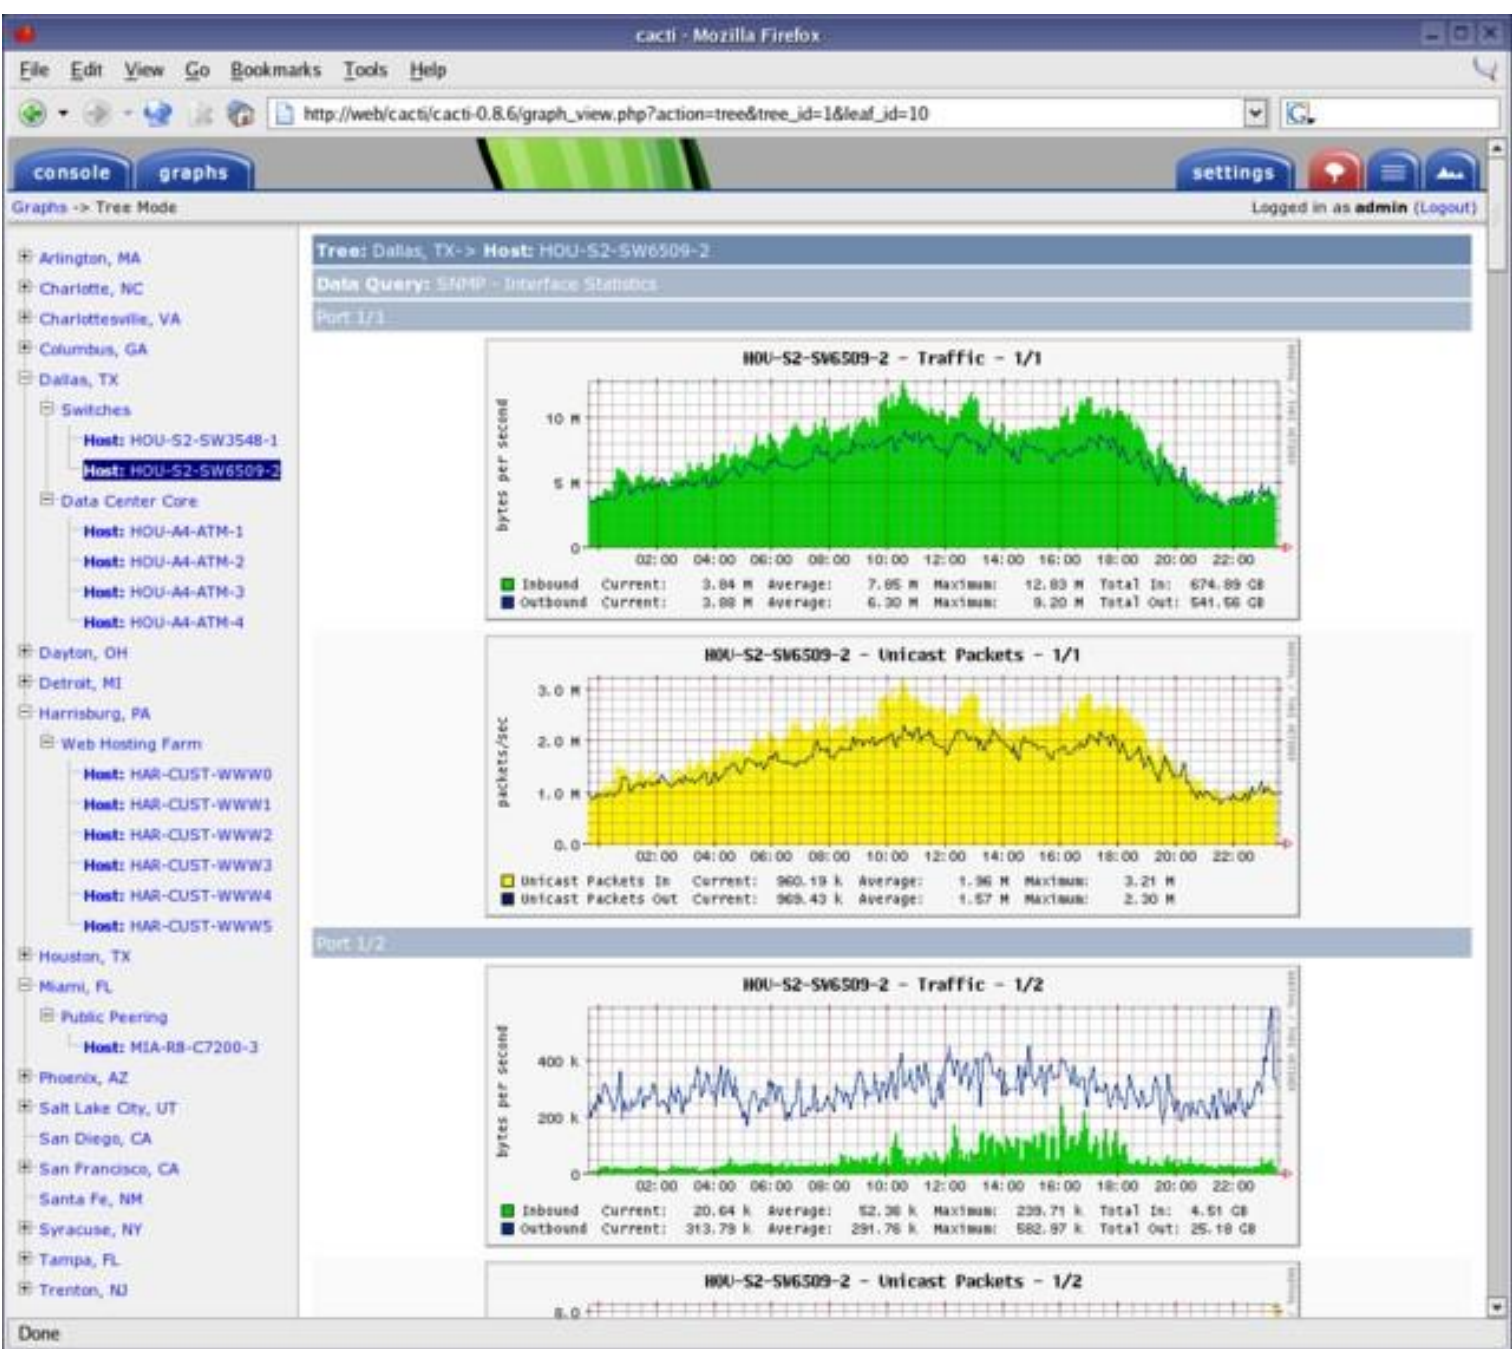


Nagios

Legacy monitoring systems

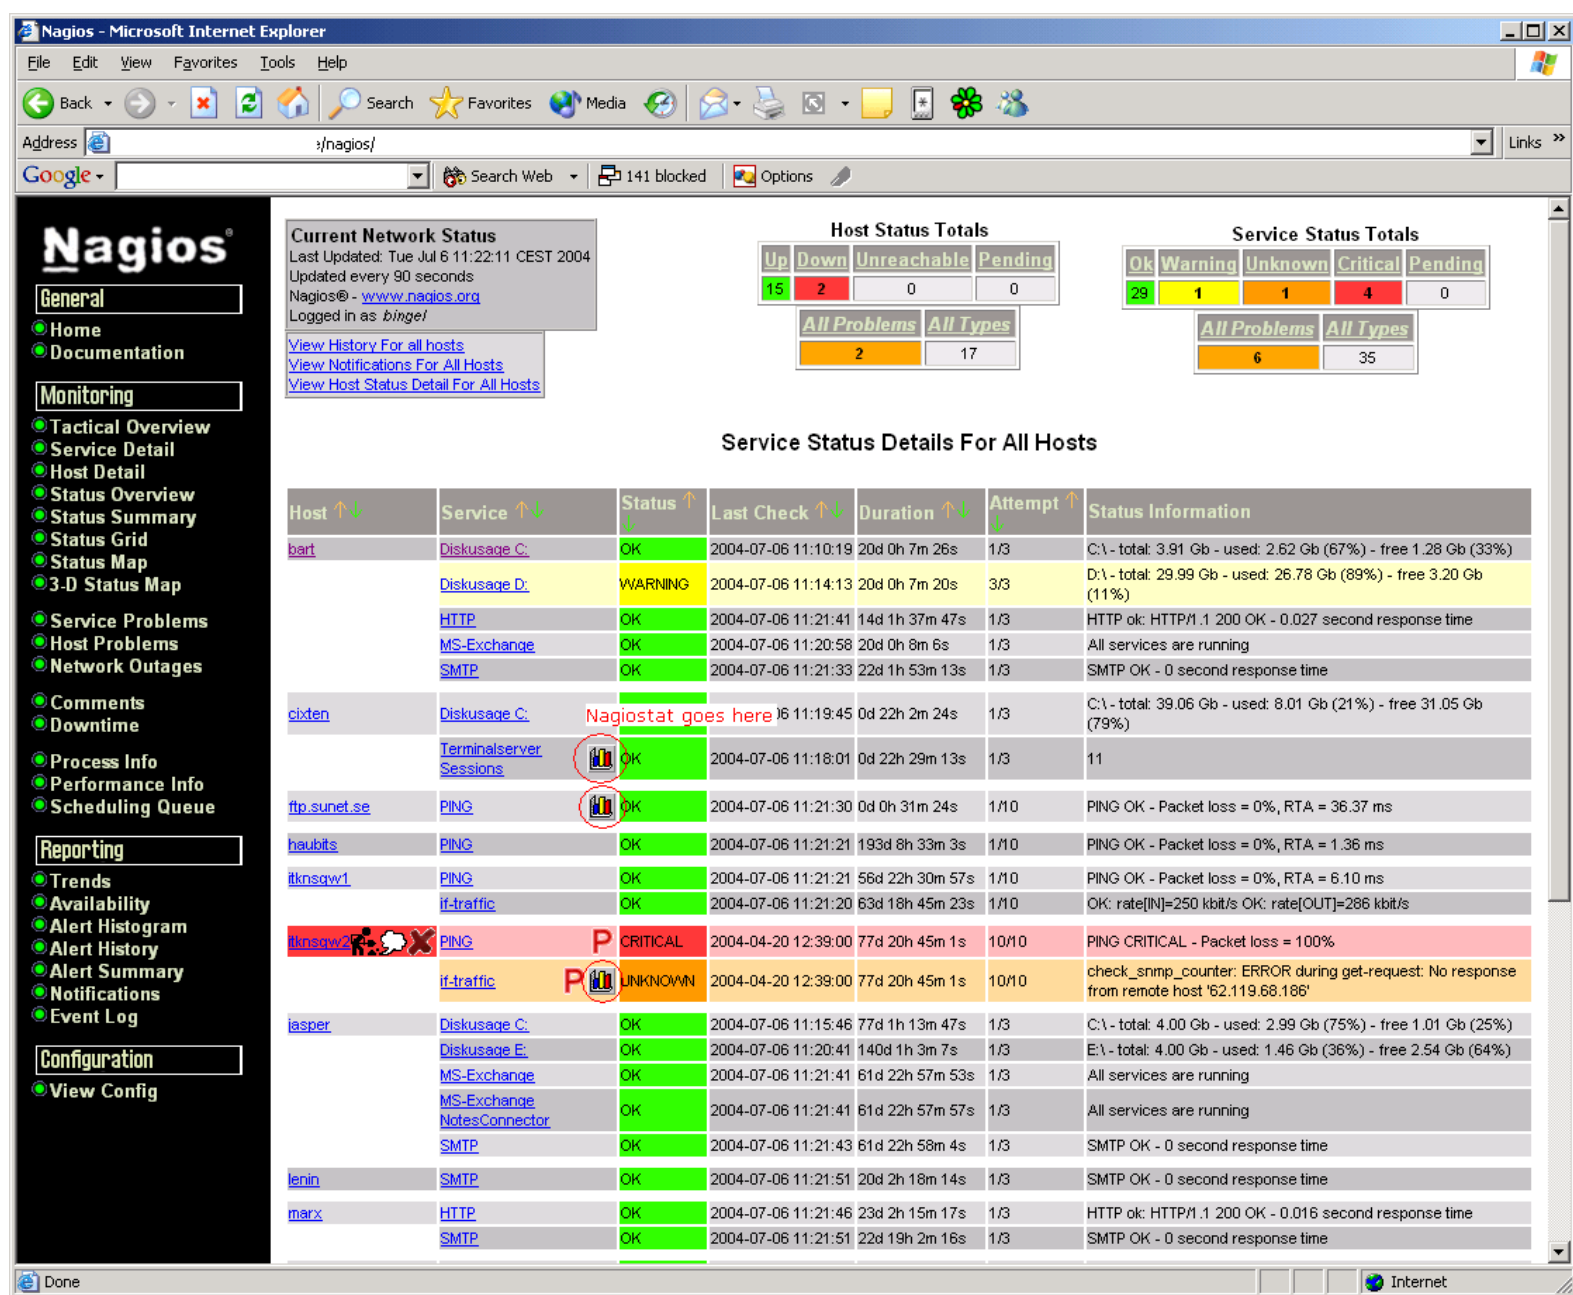


Nagios

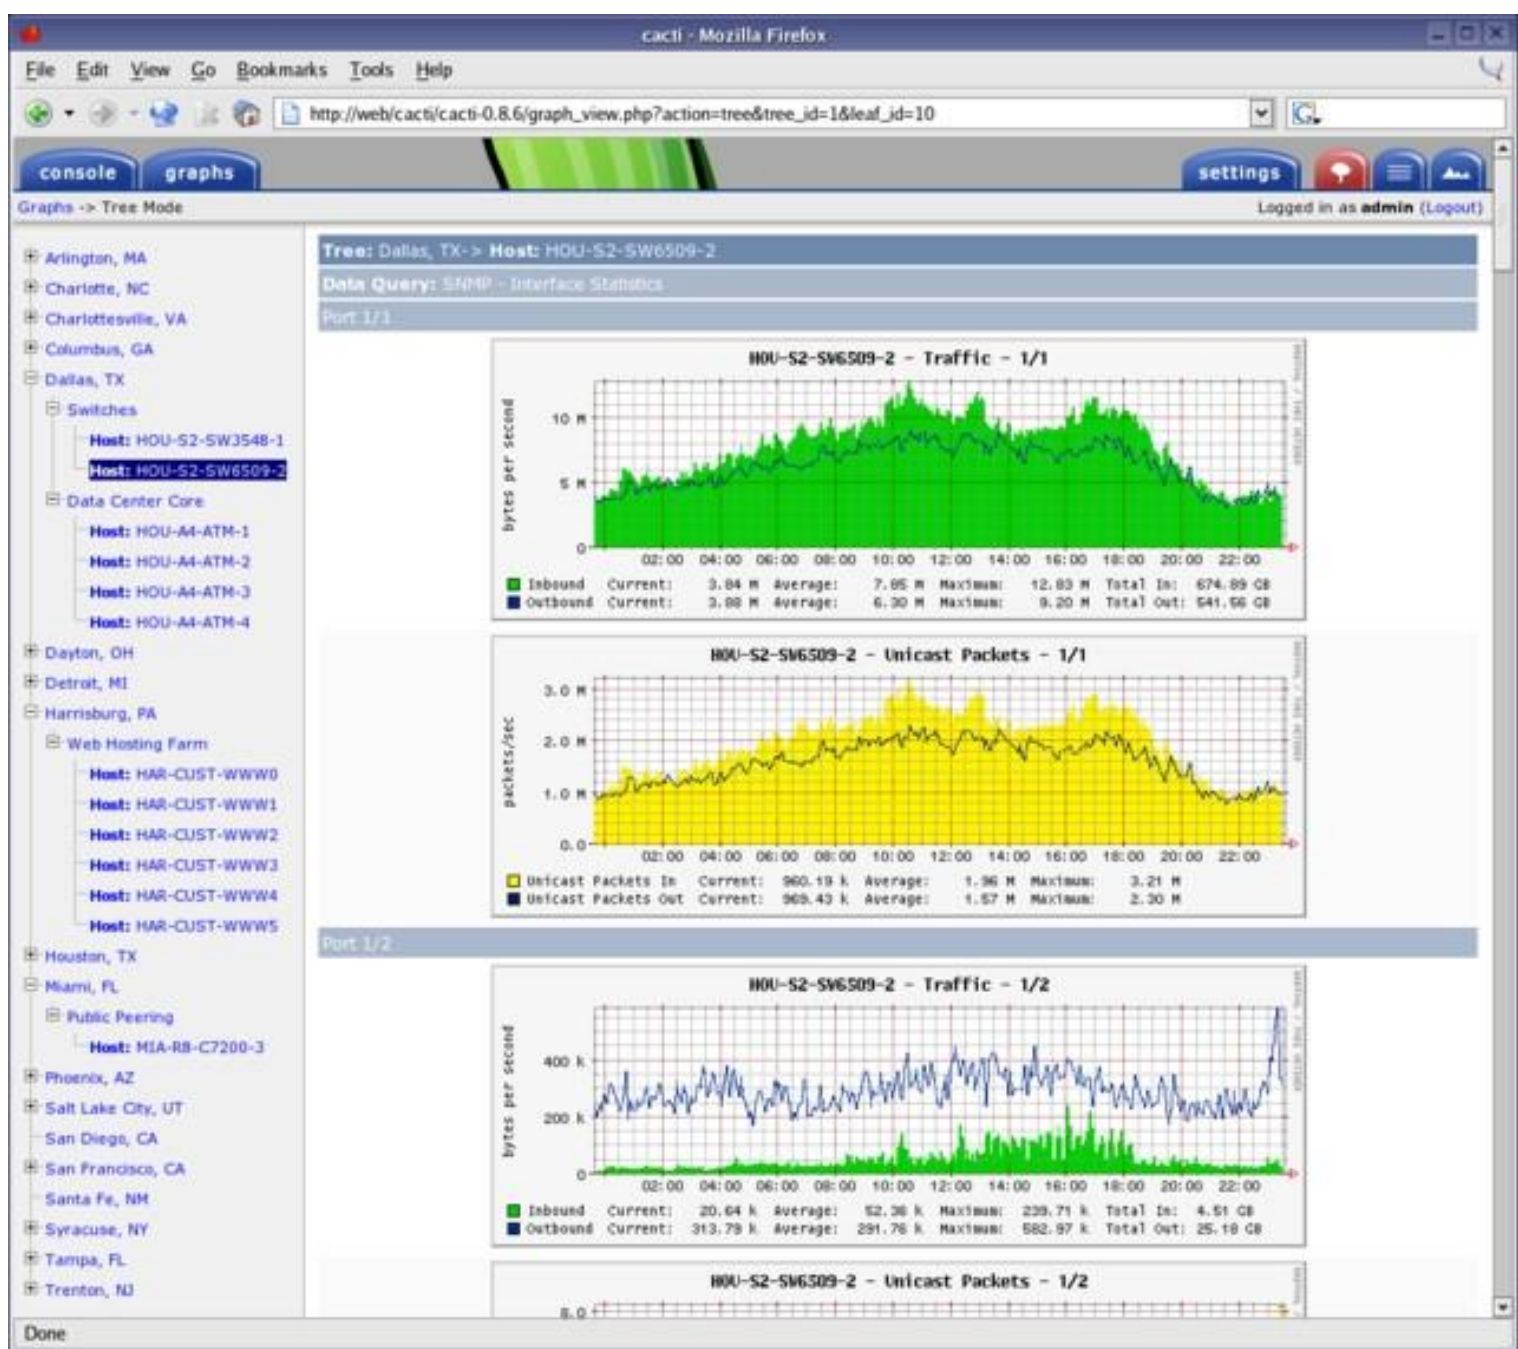


Cacti

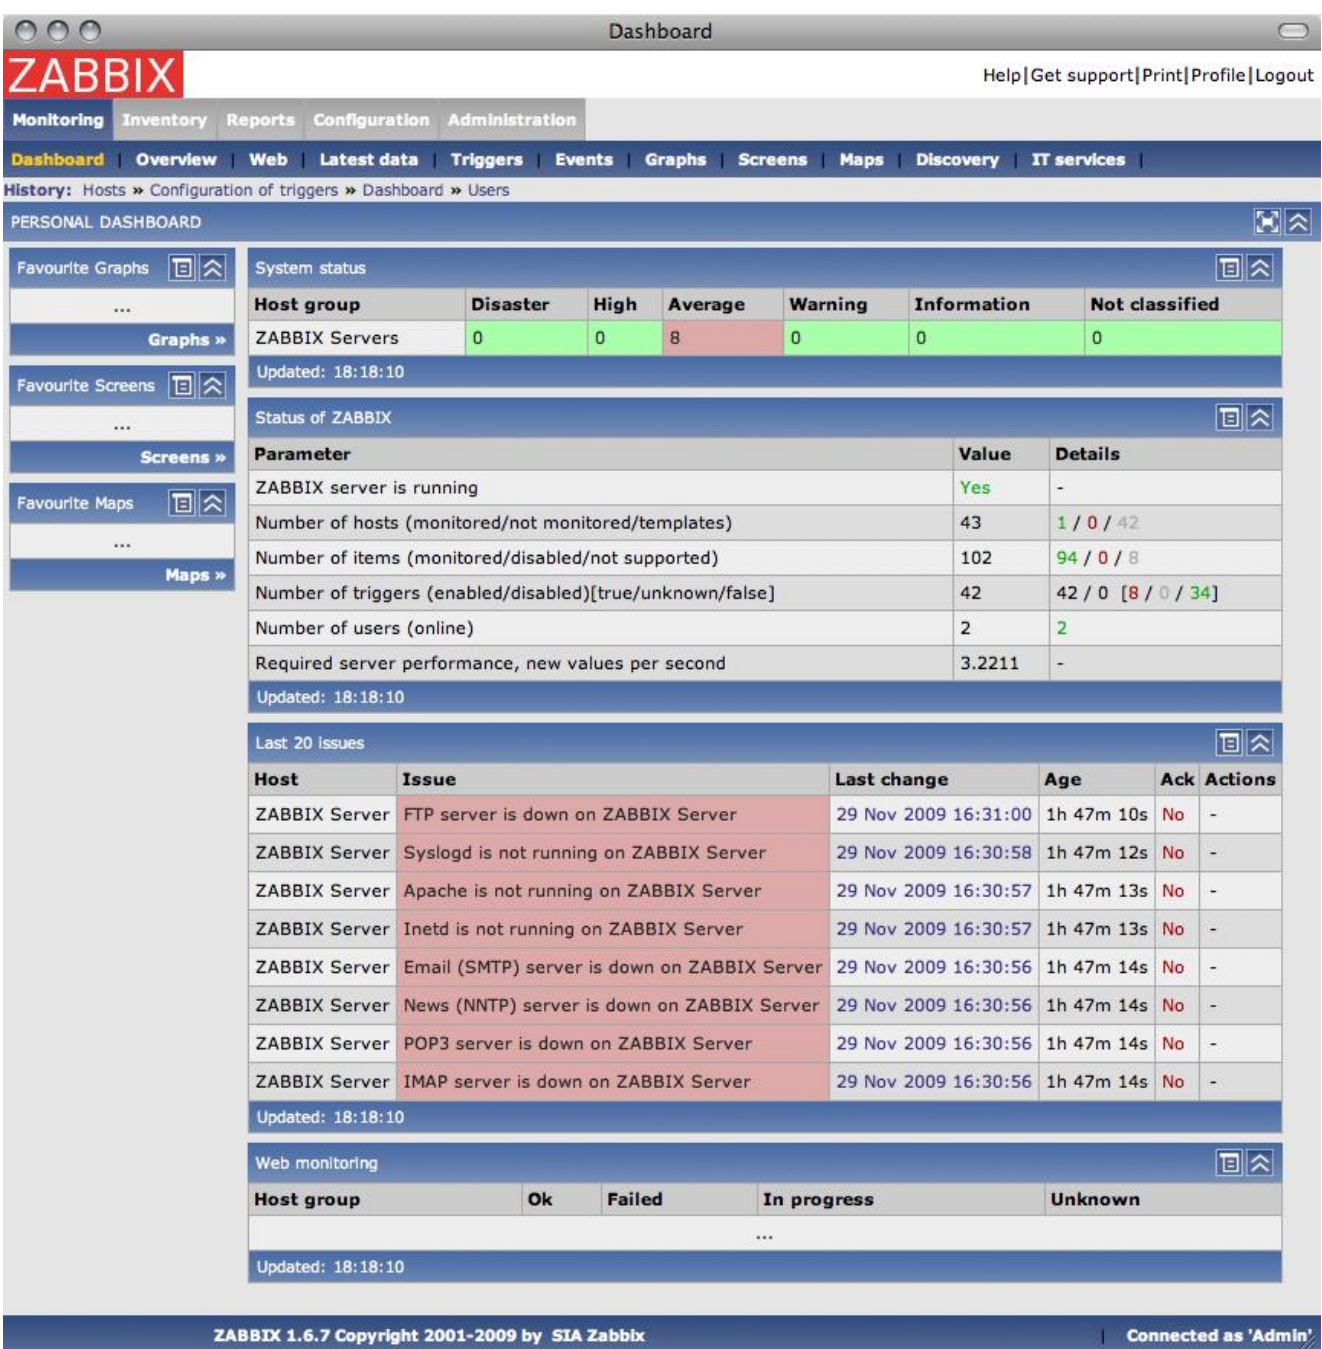
Legacy monitoring systems



Nagios



Cacti



Zabbix

Modern monitoring systems



netdata is a **monitoring agent**: you install it on all your systems:

- supports **auto-detection** and **zero configuration** for most applications and systems
- is **real-time**: every metric is on your dashboard in less than 1-second (collection to visualization)
- is **fast**: for a few thousand metrics per second, it needs < 1% CPU of a single core
- and **efficient**: it needs a few MB of RAM and no disk I/O at all while it runs
- also, it is **embeddable**, **extensible**, and **open-source** (GPL v3+)

netdata
simple. effective. awesome!
<https://my-netdata.io>

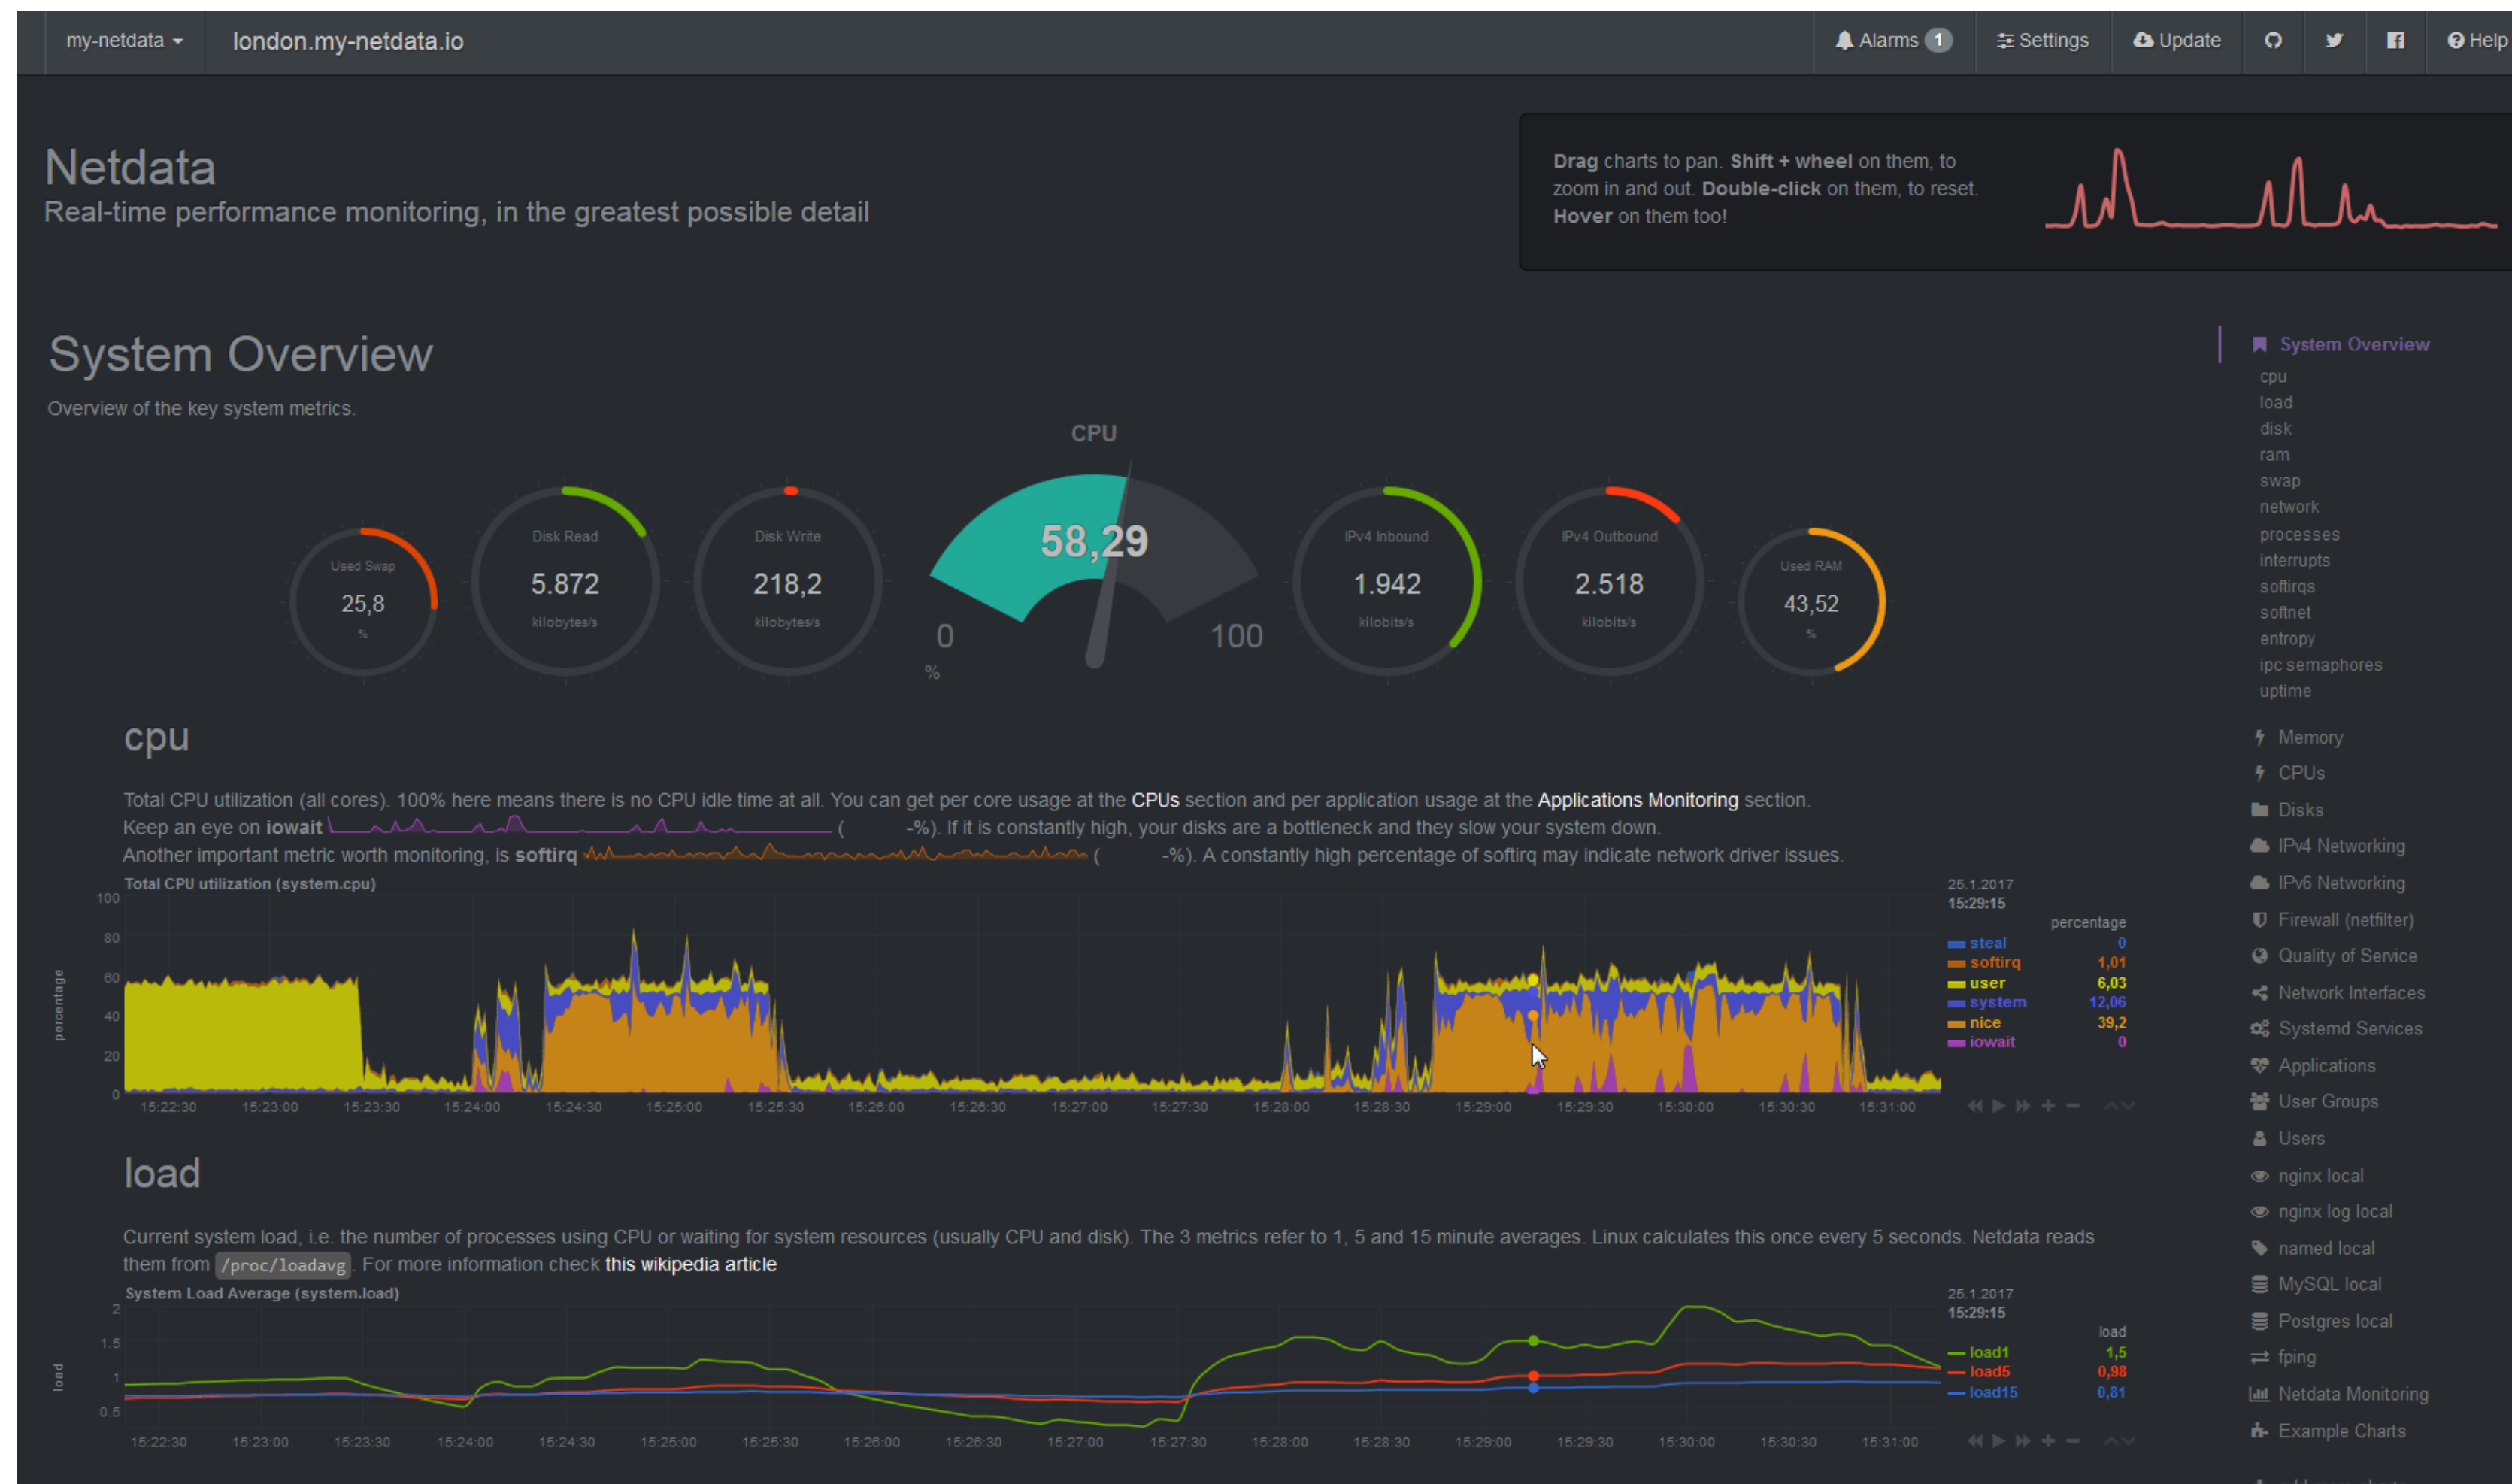
...and netdata runs everywhere:



(C) Copyright 2018
Netdata, Inc.

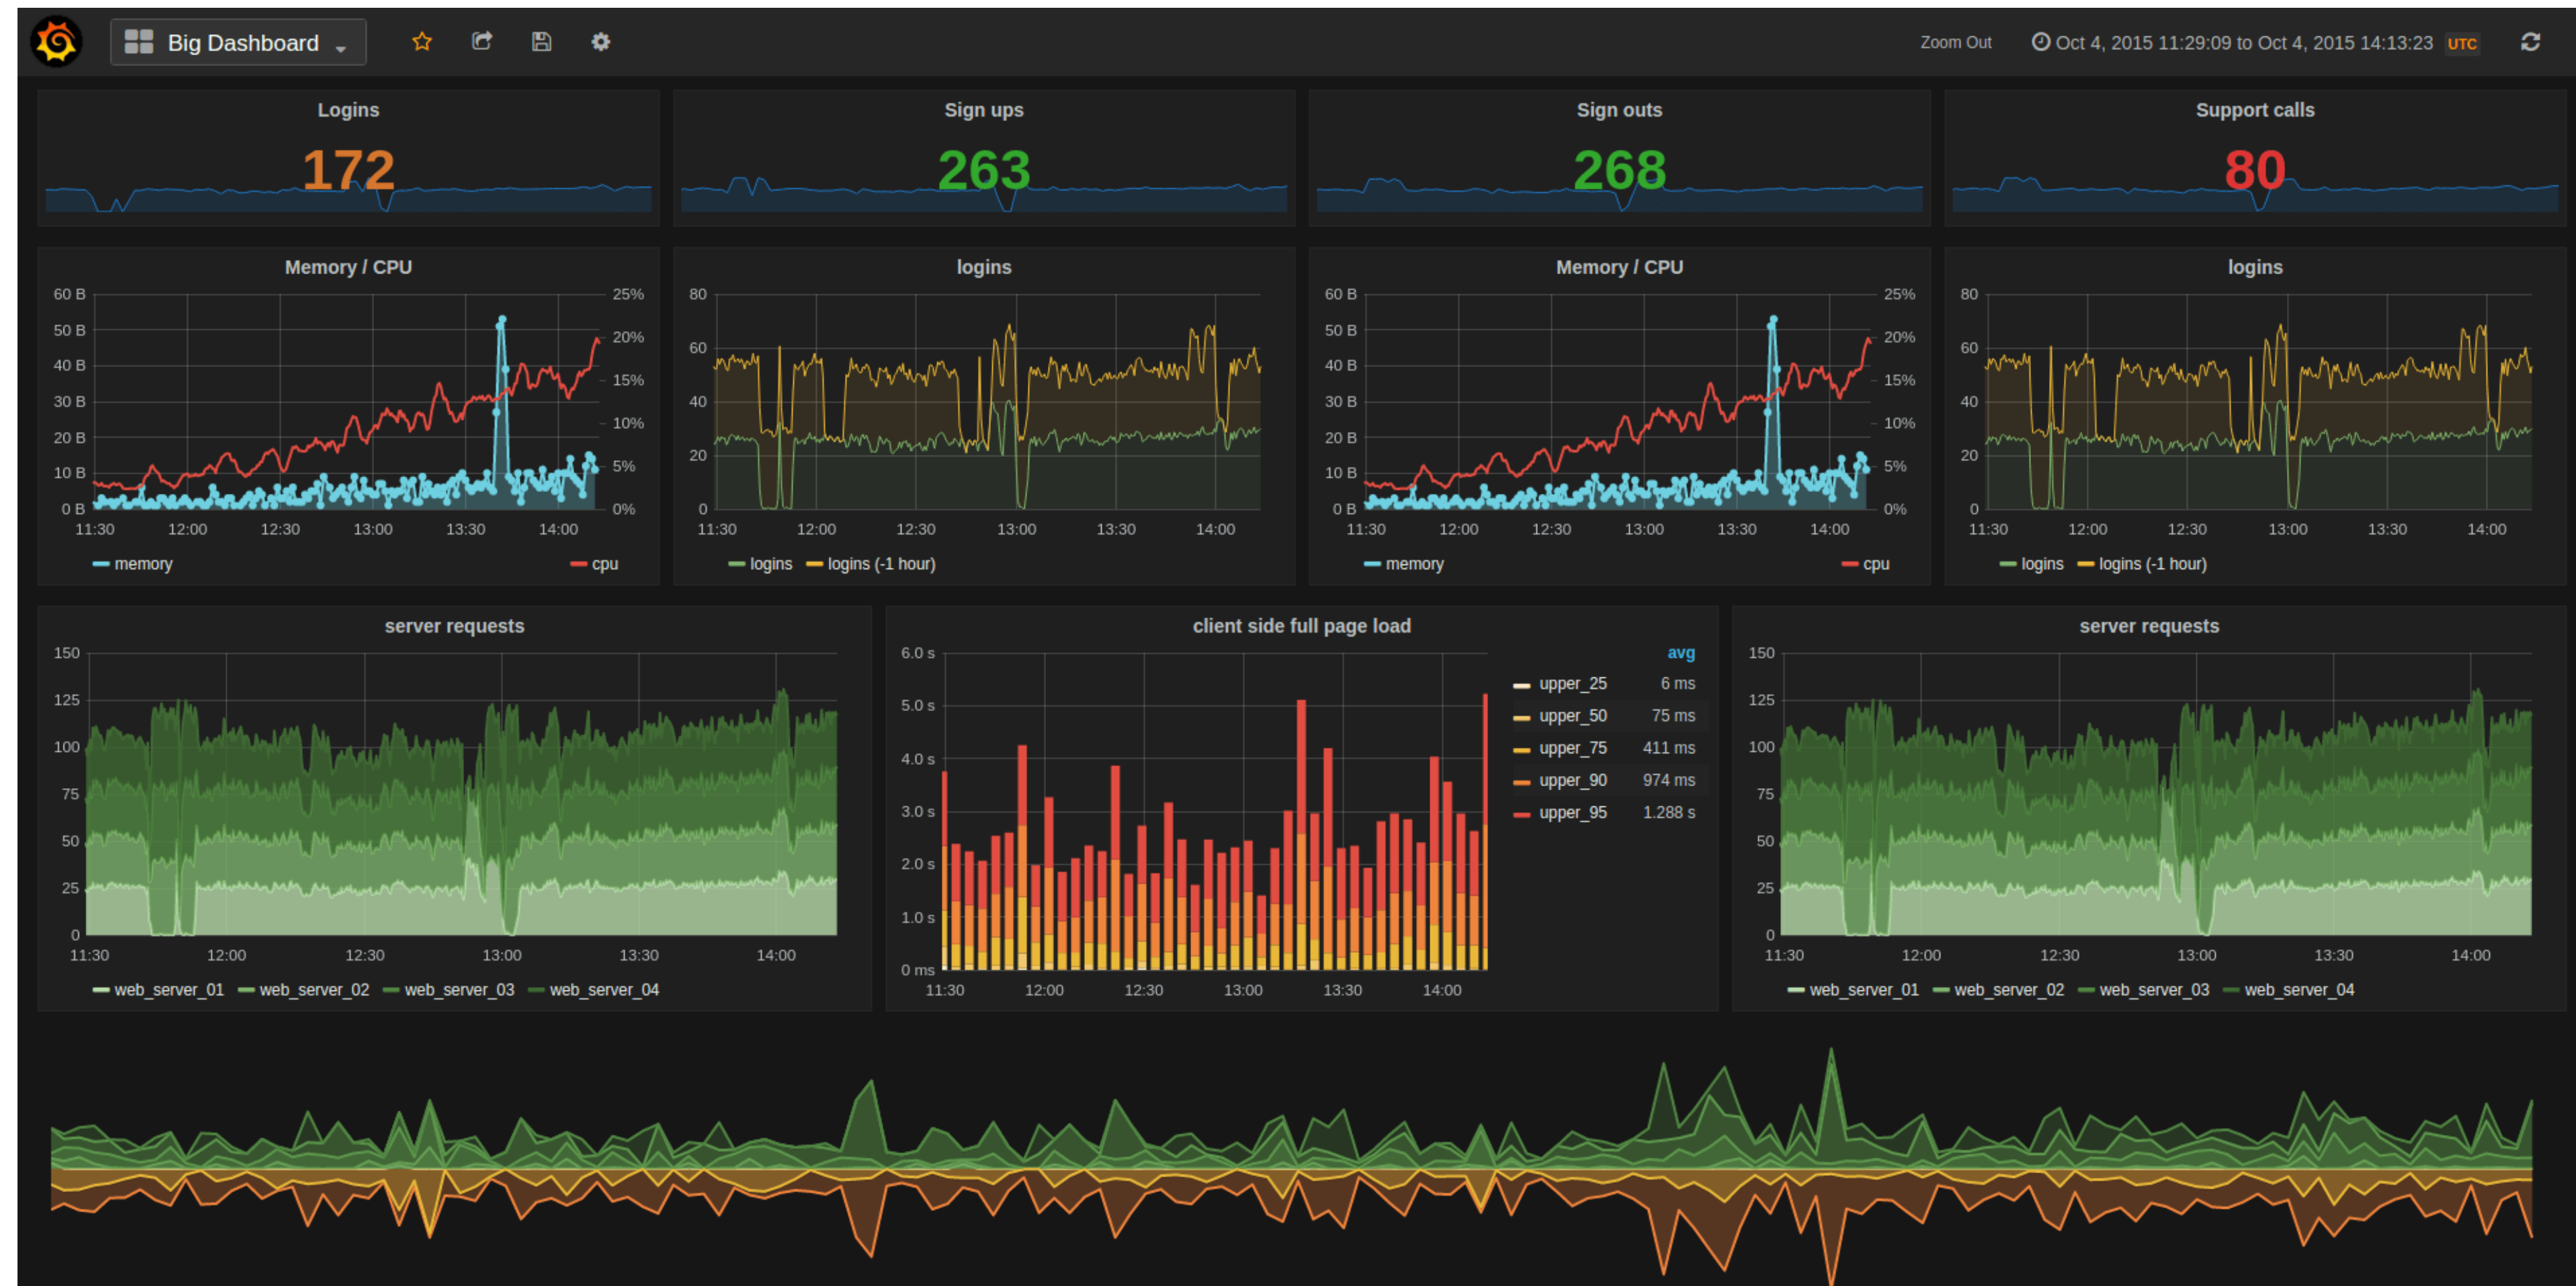
Netdata — Prometheus Exporter

Modern monitoring systems



Netdata

Modern monitoring systems



Grafana

Modern monitoring systems

The screenshot displays the Karma monitoring system interface, which is a dashboard for managing alerts. The top bar shows the number of alerts (34) and three active filters: '@cluster=HA' (1368), '@receiver=by-cluster-service' (92), and 'cluster=prod' (136). The main area is divided into several panels, each showing a different alert configuration and its history.

- Alertname: Inhibition Test Alert** (cluster: prod): Shows a history of alerts with severity: critical and severity: warning. Includes a search bar and a 'help' button.
- Alertname: Some Alerts With A Ridiculously Lo...** (cluster: prod): Shows a history of alerts with instance: server7. Includes a search bar and a 'help' button.
- Alertname: Mixed Alerts** (cluster: prod): Shows a history of alerts with instance: server6, instance: server8, and instance: server1. Includes a search bar and a 'help' button.
- Alertname: Disk Free Low** (cluster: prod): Shows a history of alerts with summary: Only 4% free space left on /disk, device: /dev/sda0, instance: server0, and other disk-related alerts. Includes a search bar and a 'help' button.
- Alertname: Always Silenced Alert** (cluster: prod): Shows a history of alerts with instance: server1, job: mysql_exporter, region: SA, severity: info, @cluster: HA, @receiver: by-cluster-service, and alertReference. Includes a search bar and a 'help' button.
- Alertname: Time Annotation** (cluster: prod): Shows a history of alerts with instance: server1, job: ntp_exporter, region: AP, severity: warning, @cluster: HA, @receiver: by-cluster-service. Includes a search bar and a 'help' button.
- Alertname: Always On Alert** (cluster: prod): Shows a history of alerts with instance: server1 and instance: server2. Includes a search bar and a 'help' button.

Each panel also includes a 'Silence' button and a 'help' button. The interface is designed to be highly customizable and informative, providing a clear overview of the system's health and any issues that arise.

Karma

Legacy monitoring systems



Now we are here

- ~~• The way from Legacy monitoring systems to Modern~~
- **Monitoring Problems**
- Legacy Infrastructure Service Discovery
- Prometheus Operator as the Solution
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

Monitoring Problems

- Ops is the bottleneck

Monitoring Problems

- Ops is the bottleneck
- Time to market of business metrics

Monitoring Problems

- Ops is the bottleneck
- Time to market of business metrics
- Metrics, Dashboards, Alerts As A Code

Monitoring Problems

- Ops is the bottleneck
- Time to market of business metrics
- Metrics, Dashboards, Alerts As A Code
- More then one monitoring system

Now we are here


- ~~• The way from Legacy monitoring systems to Modern~~
- ~~• Monitoring Problems~~
- Legacy Infrastructure Service Discovery
- Prometheus Operator as the Solution
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

Prometheus Service Discovery with Consul




Prometheus

Metrics




HashiCorp
Consul

Discovery



NETDATA

Agent



Prometheus Service Discovery with Consul

```
1  ...  additionalScrapeConfigs:
2  ...  - job_name: consul
3  ...    consul_sd_configs:
4  ...      - server: 'consul-server.consul:8500'
5  ...        datacenter: '{{ consul_dc }}'
6  ...        tags: ["monitoring"]
7  ...    relabel_configs:
8  ...      - source_labels: [__meta_consul_service]
9  ...        target_label: job
10 ...      - source_labels: [__meta_consul_node]
11 ...        target_label: instance
12
13 ...  - job_name: netdata
14 ...    metrics_path: '/api/v1/allmetrics'
15 ...    params:
16 ...      format: [prometheus]
17 ...    consul_sd_configs:
18 ...      - server: 'consul-server.consul:8500'
19 ...        datacenter: '{{ consul_dc }}'
20 ...        tags: ["netdata"]
21 ...    relabel_configs:
22 ...      - source_labels: [__meta_consul_service]
23 ...        target_label: job
24 ...      - source_labels: [__meta_consul_node]
25 ...        target_label: instance
```


Now we are here

- ~~• The way from Legacy monitoring systems to Modern~~
- ~~• Monitoring Problems~~
- ~~• Legacy Infrastructure Service Discovery~~
- **Prometheus Operator as the Solution**
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

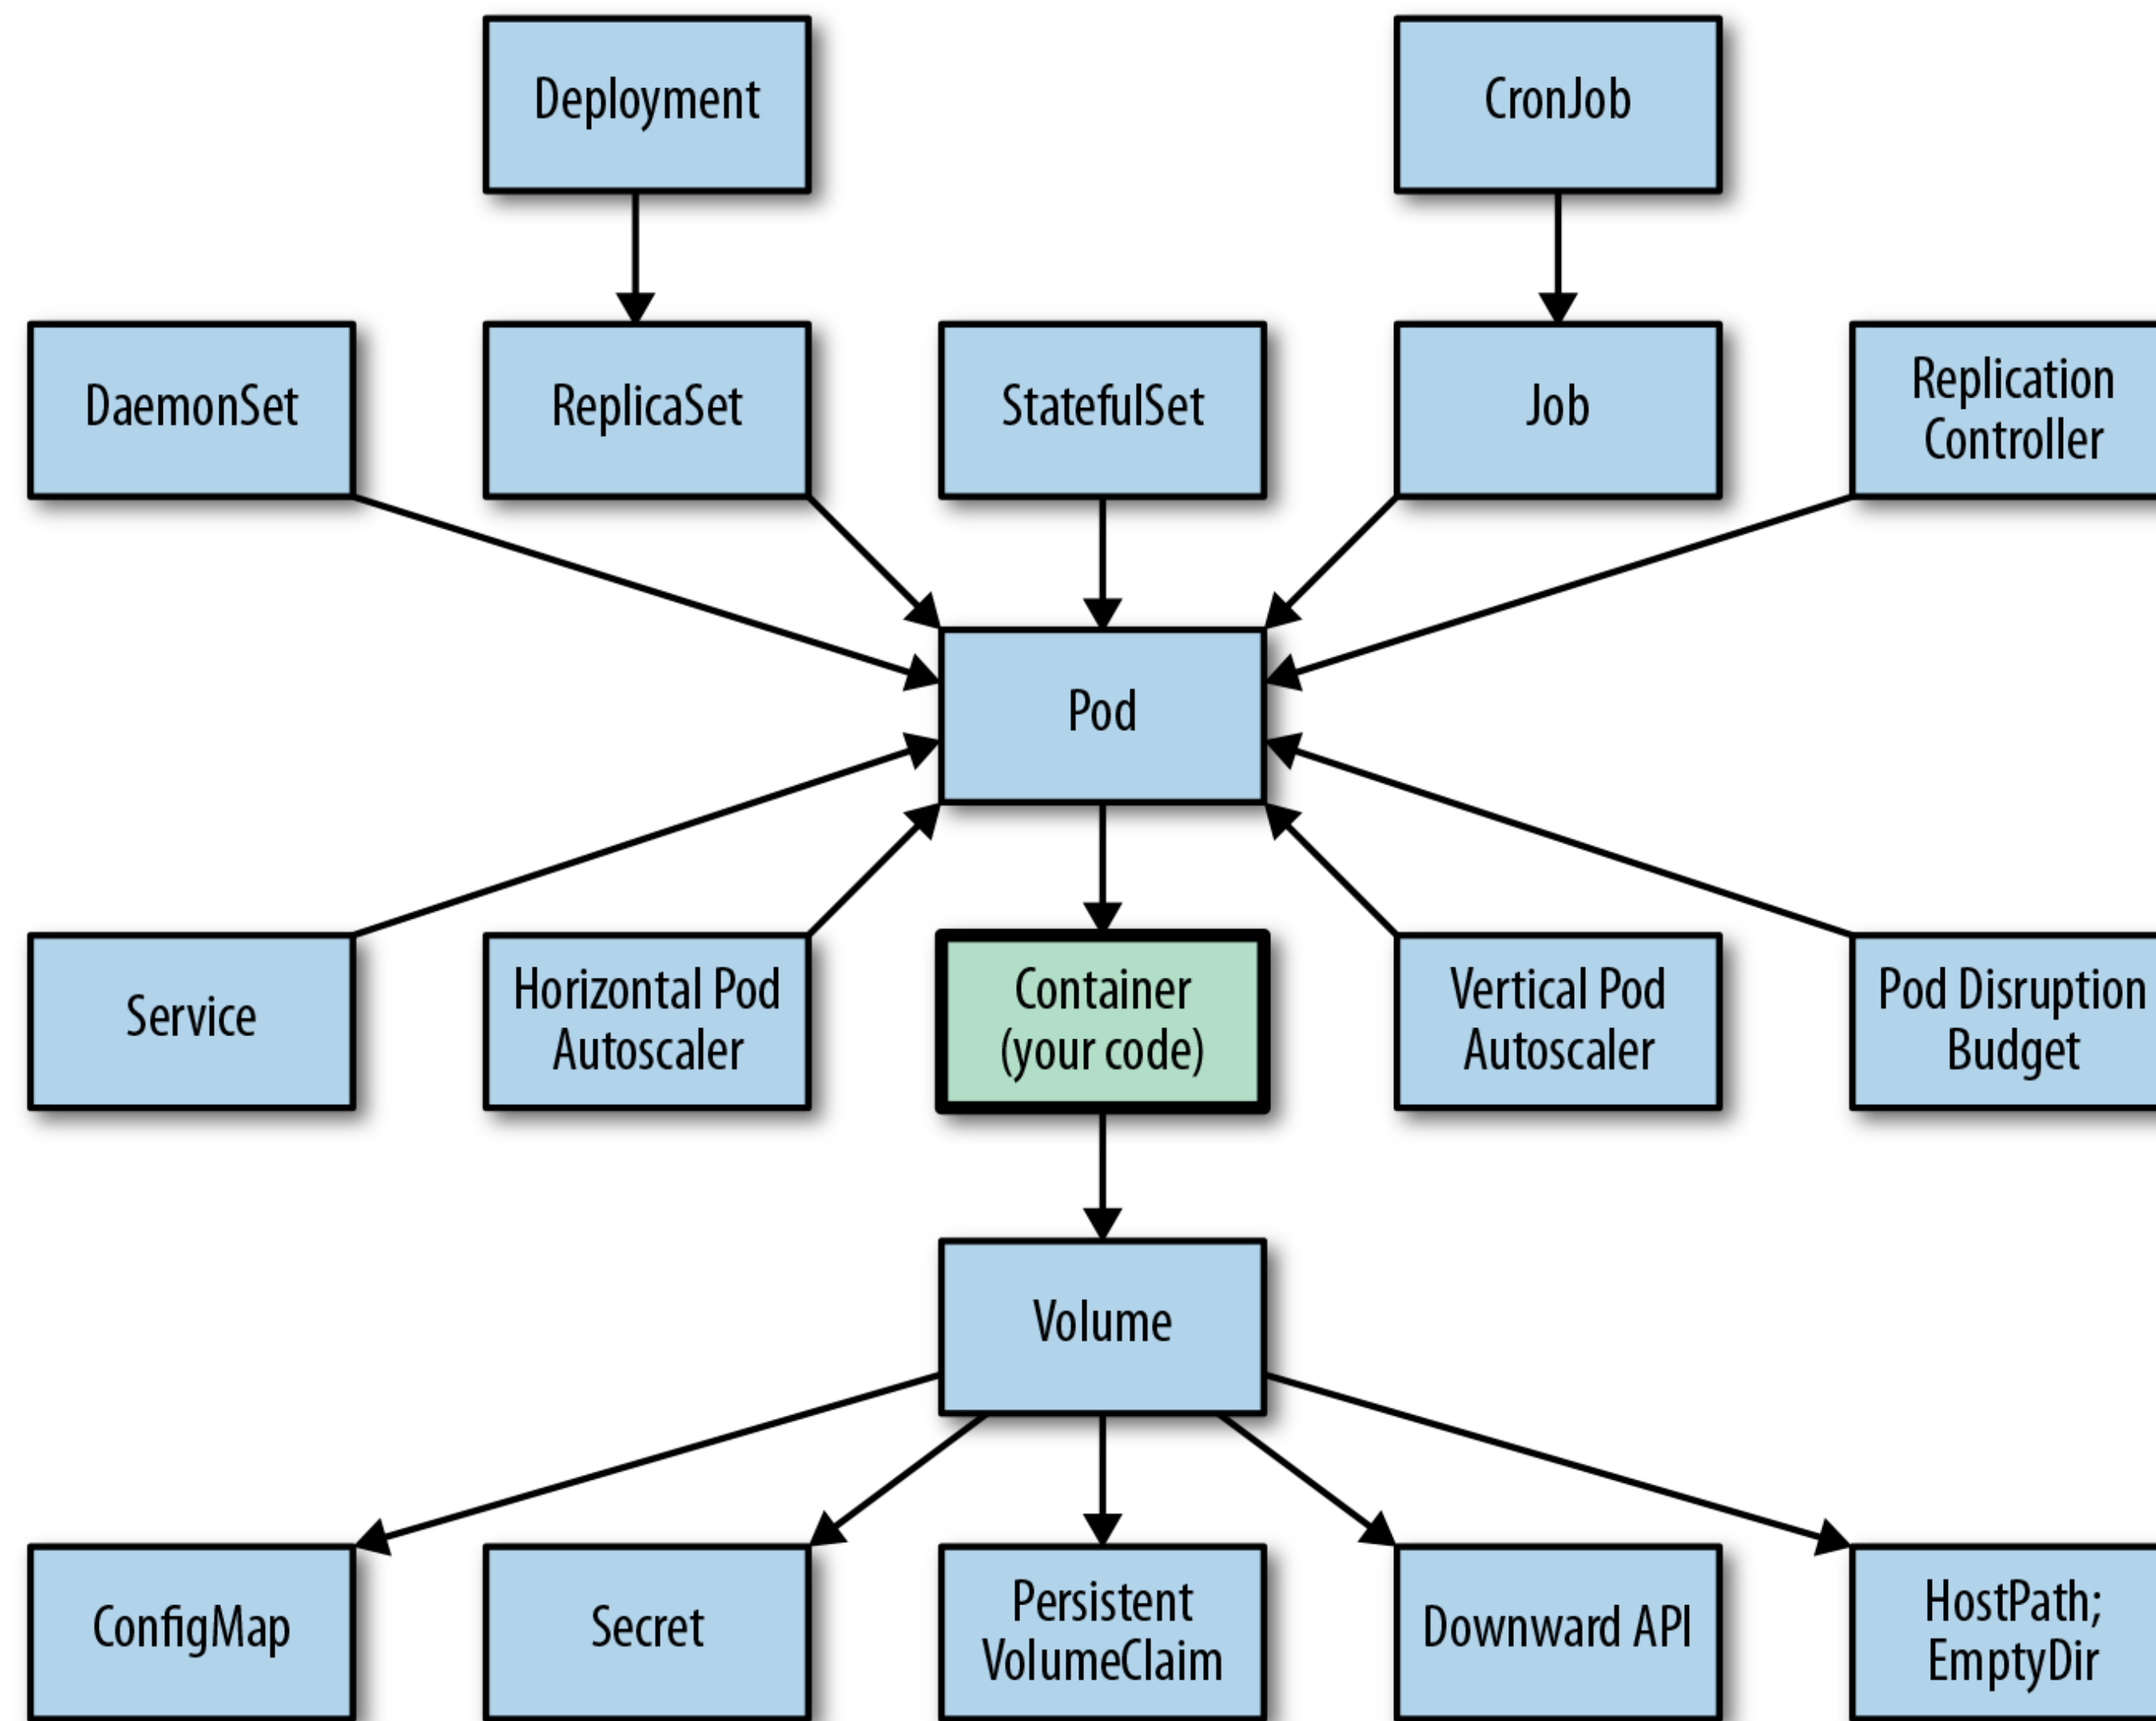
Kubernetes




WHAT GIVES PEOPLE
FEELINGS OF POWER



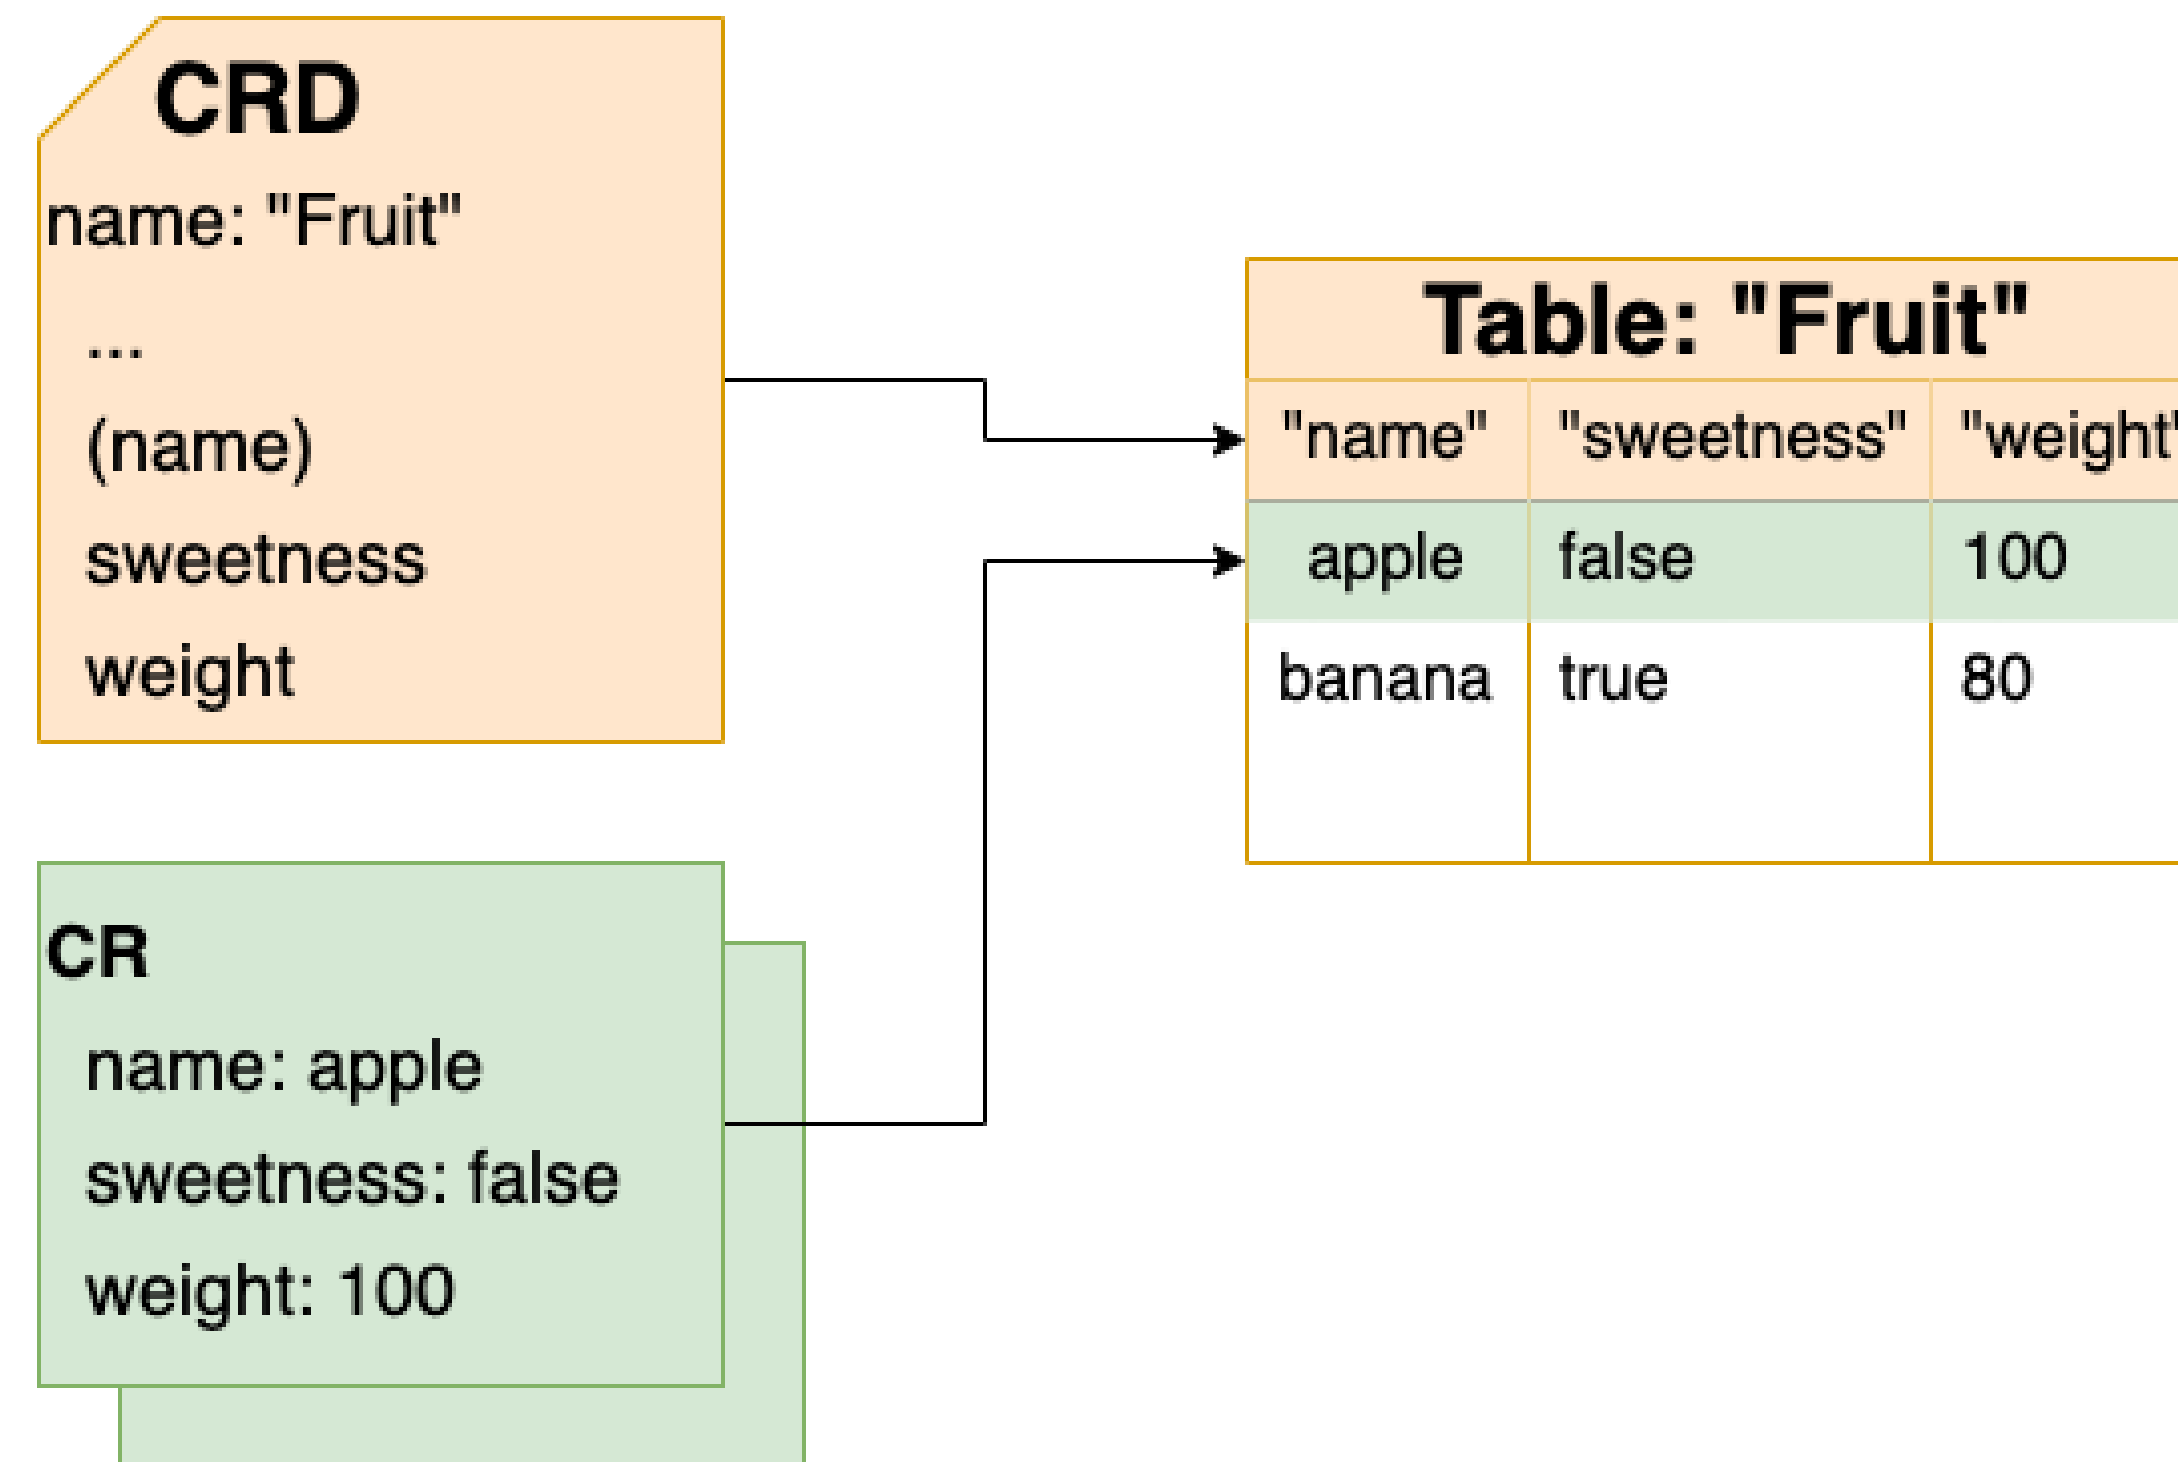
Kubernetes



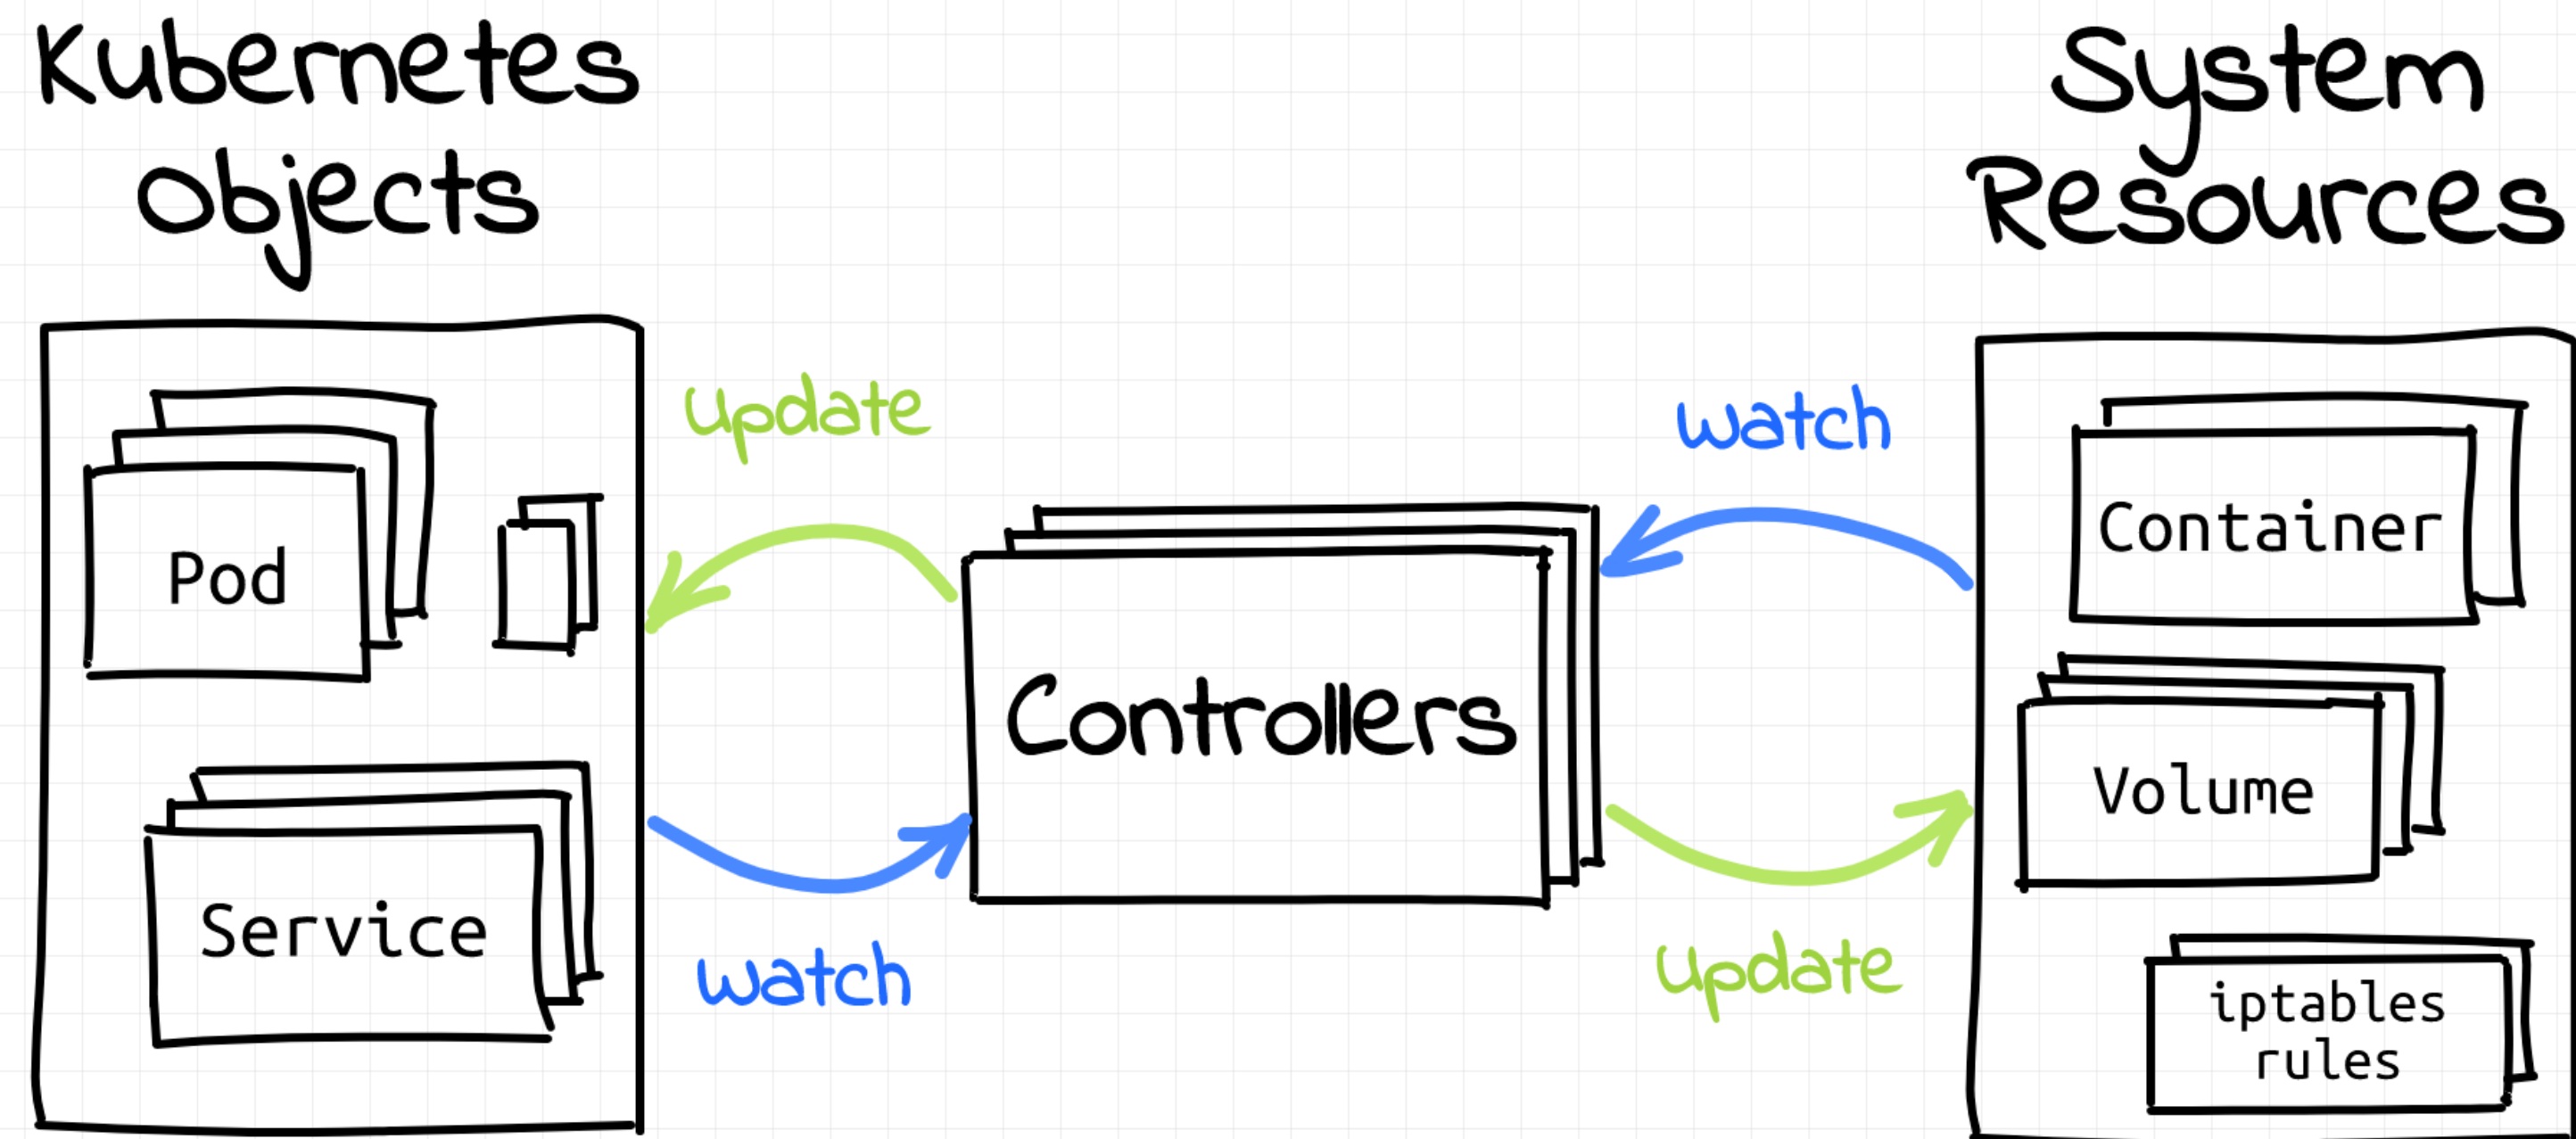
Kubernetes Custom Resource Definition



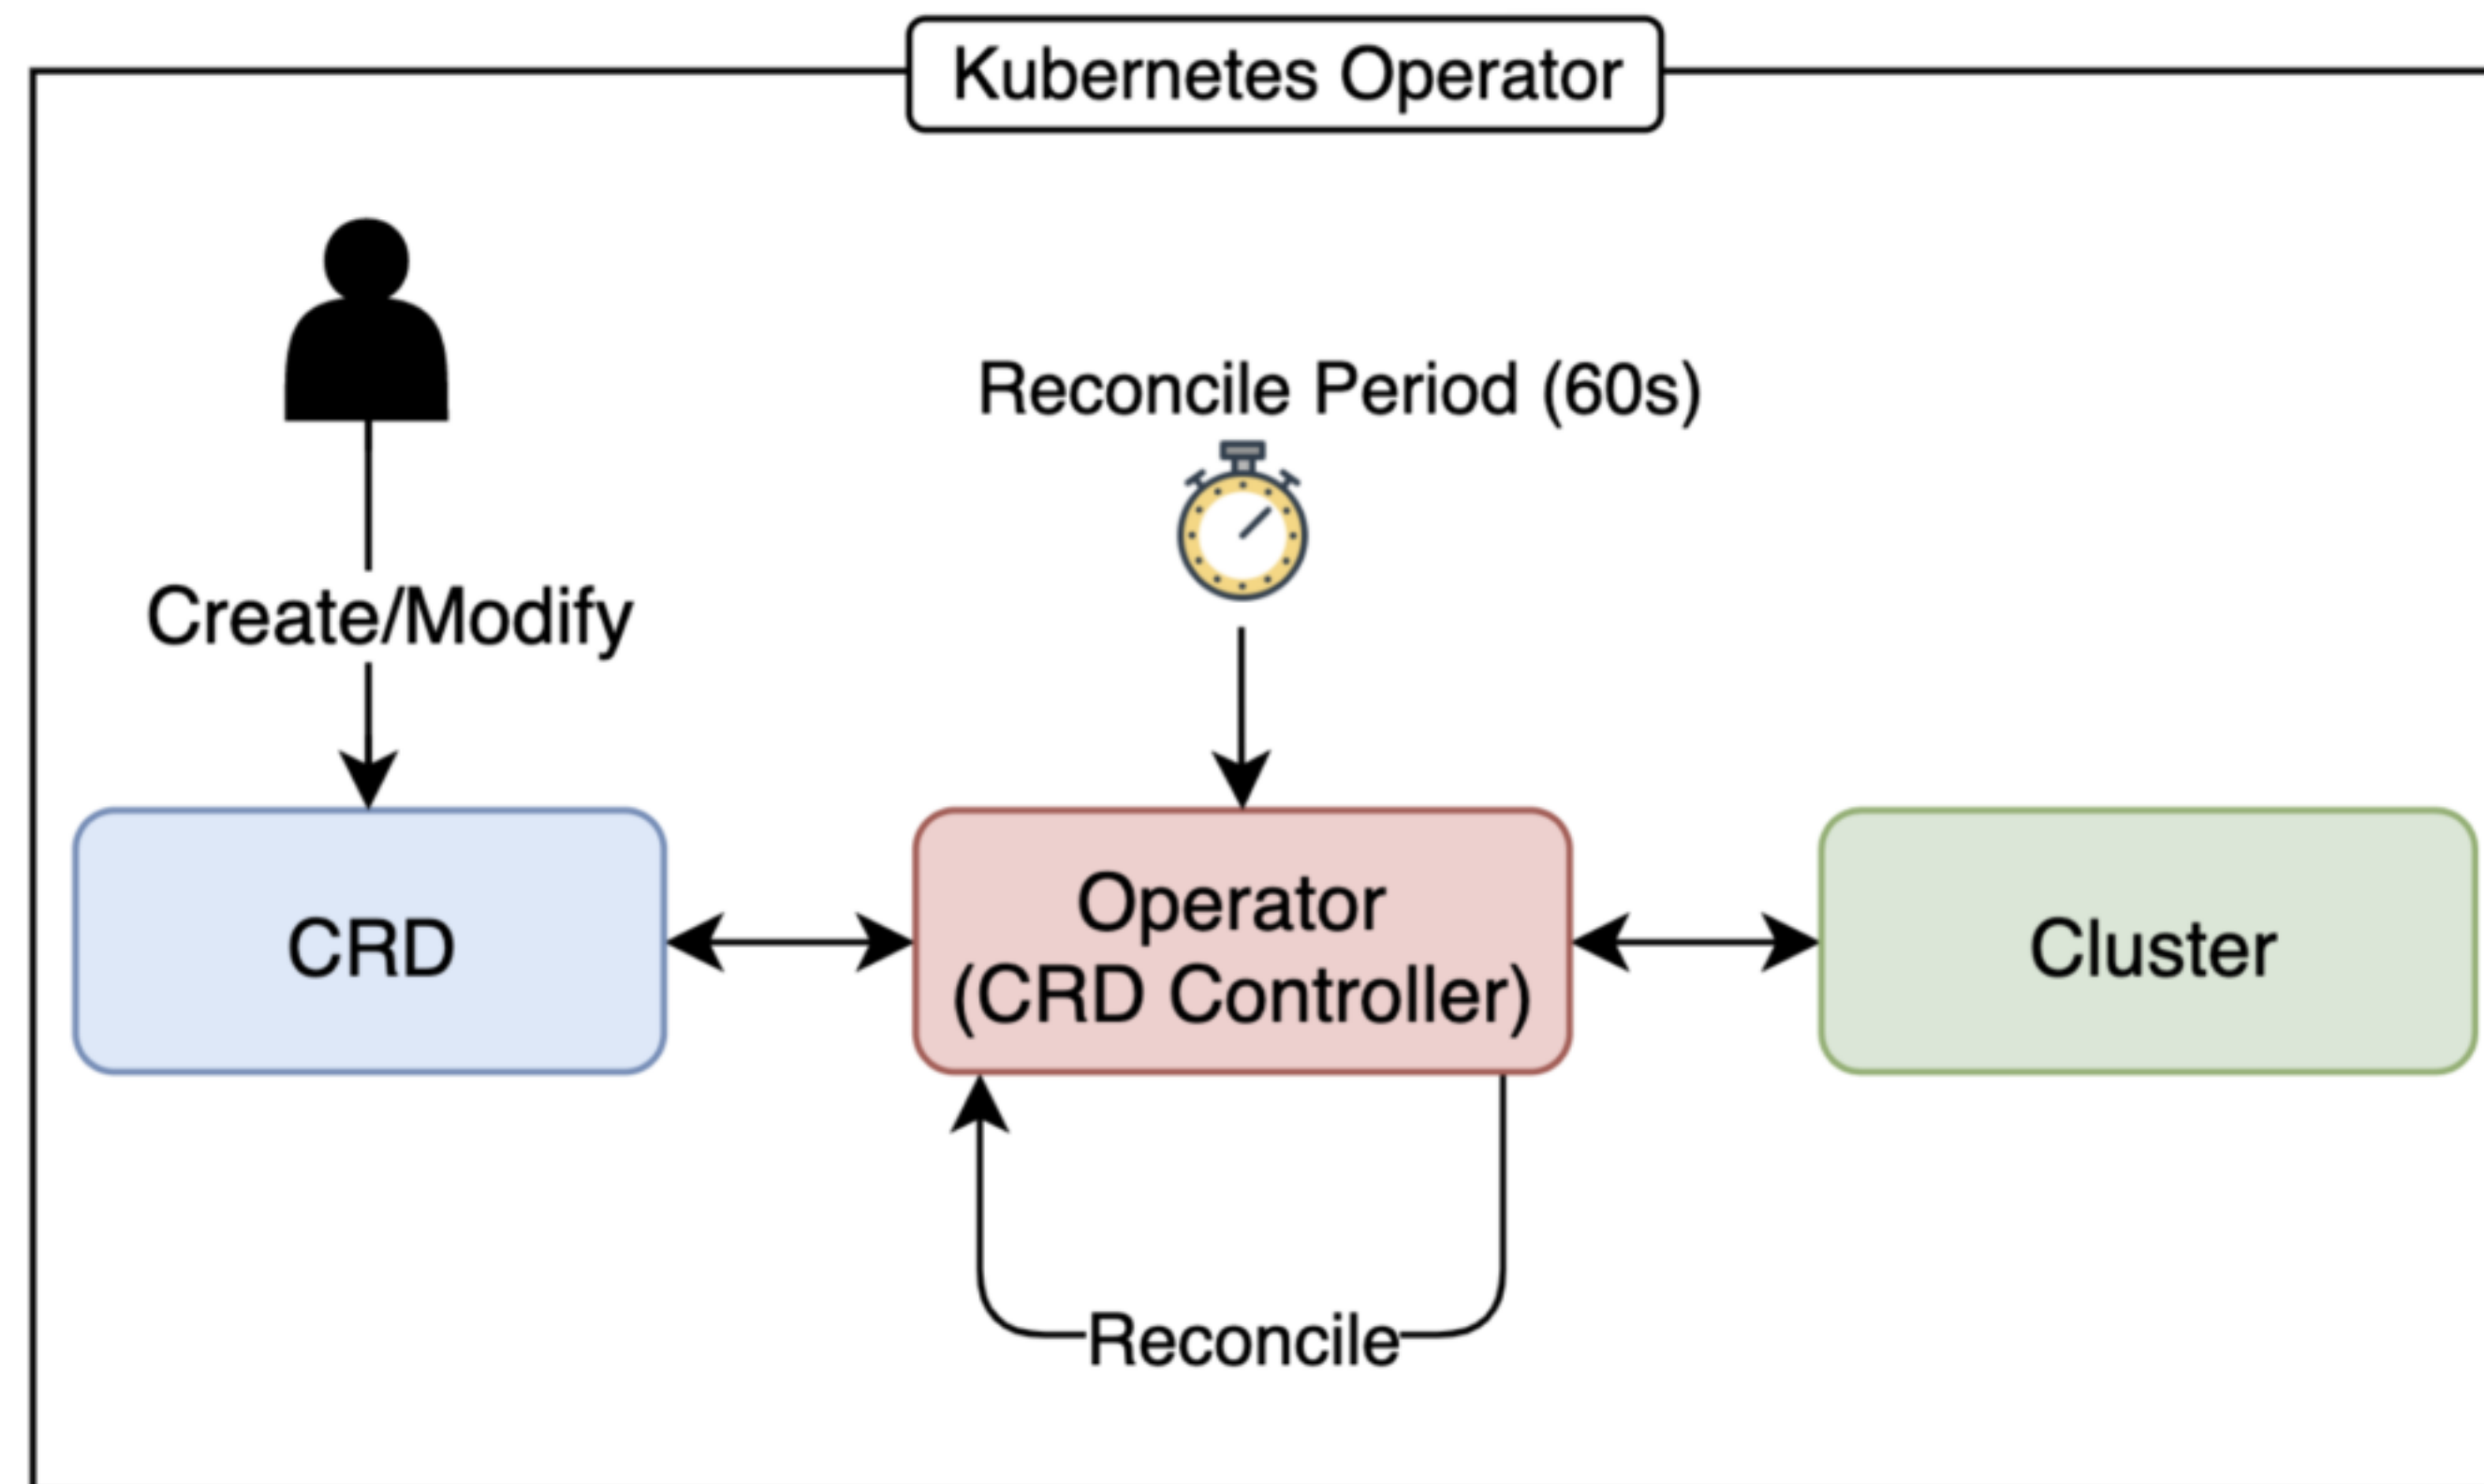
Kubernetes Custom Resource Definition



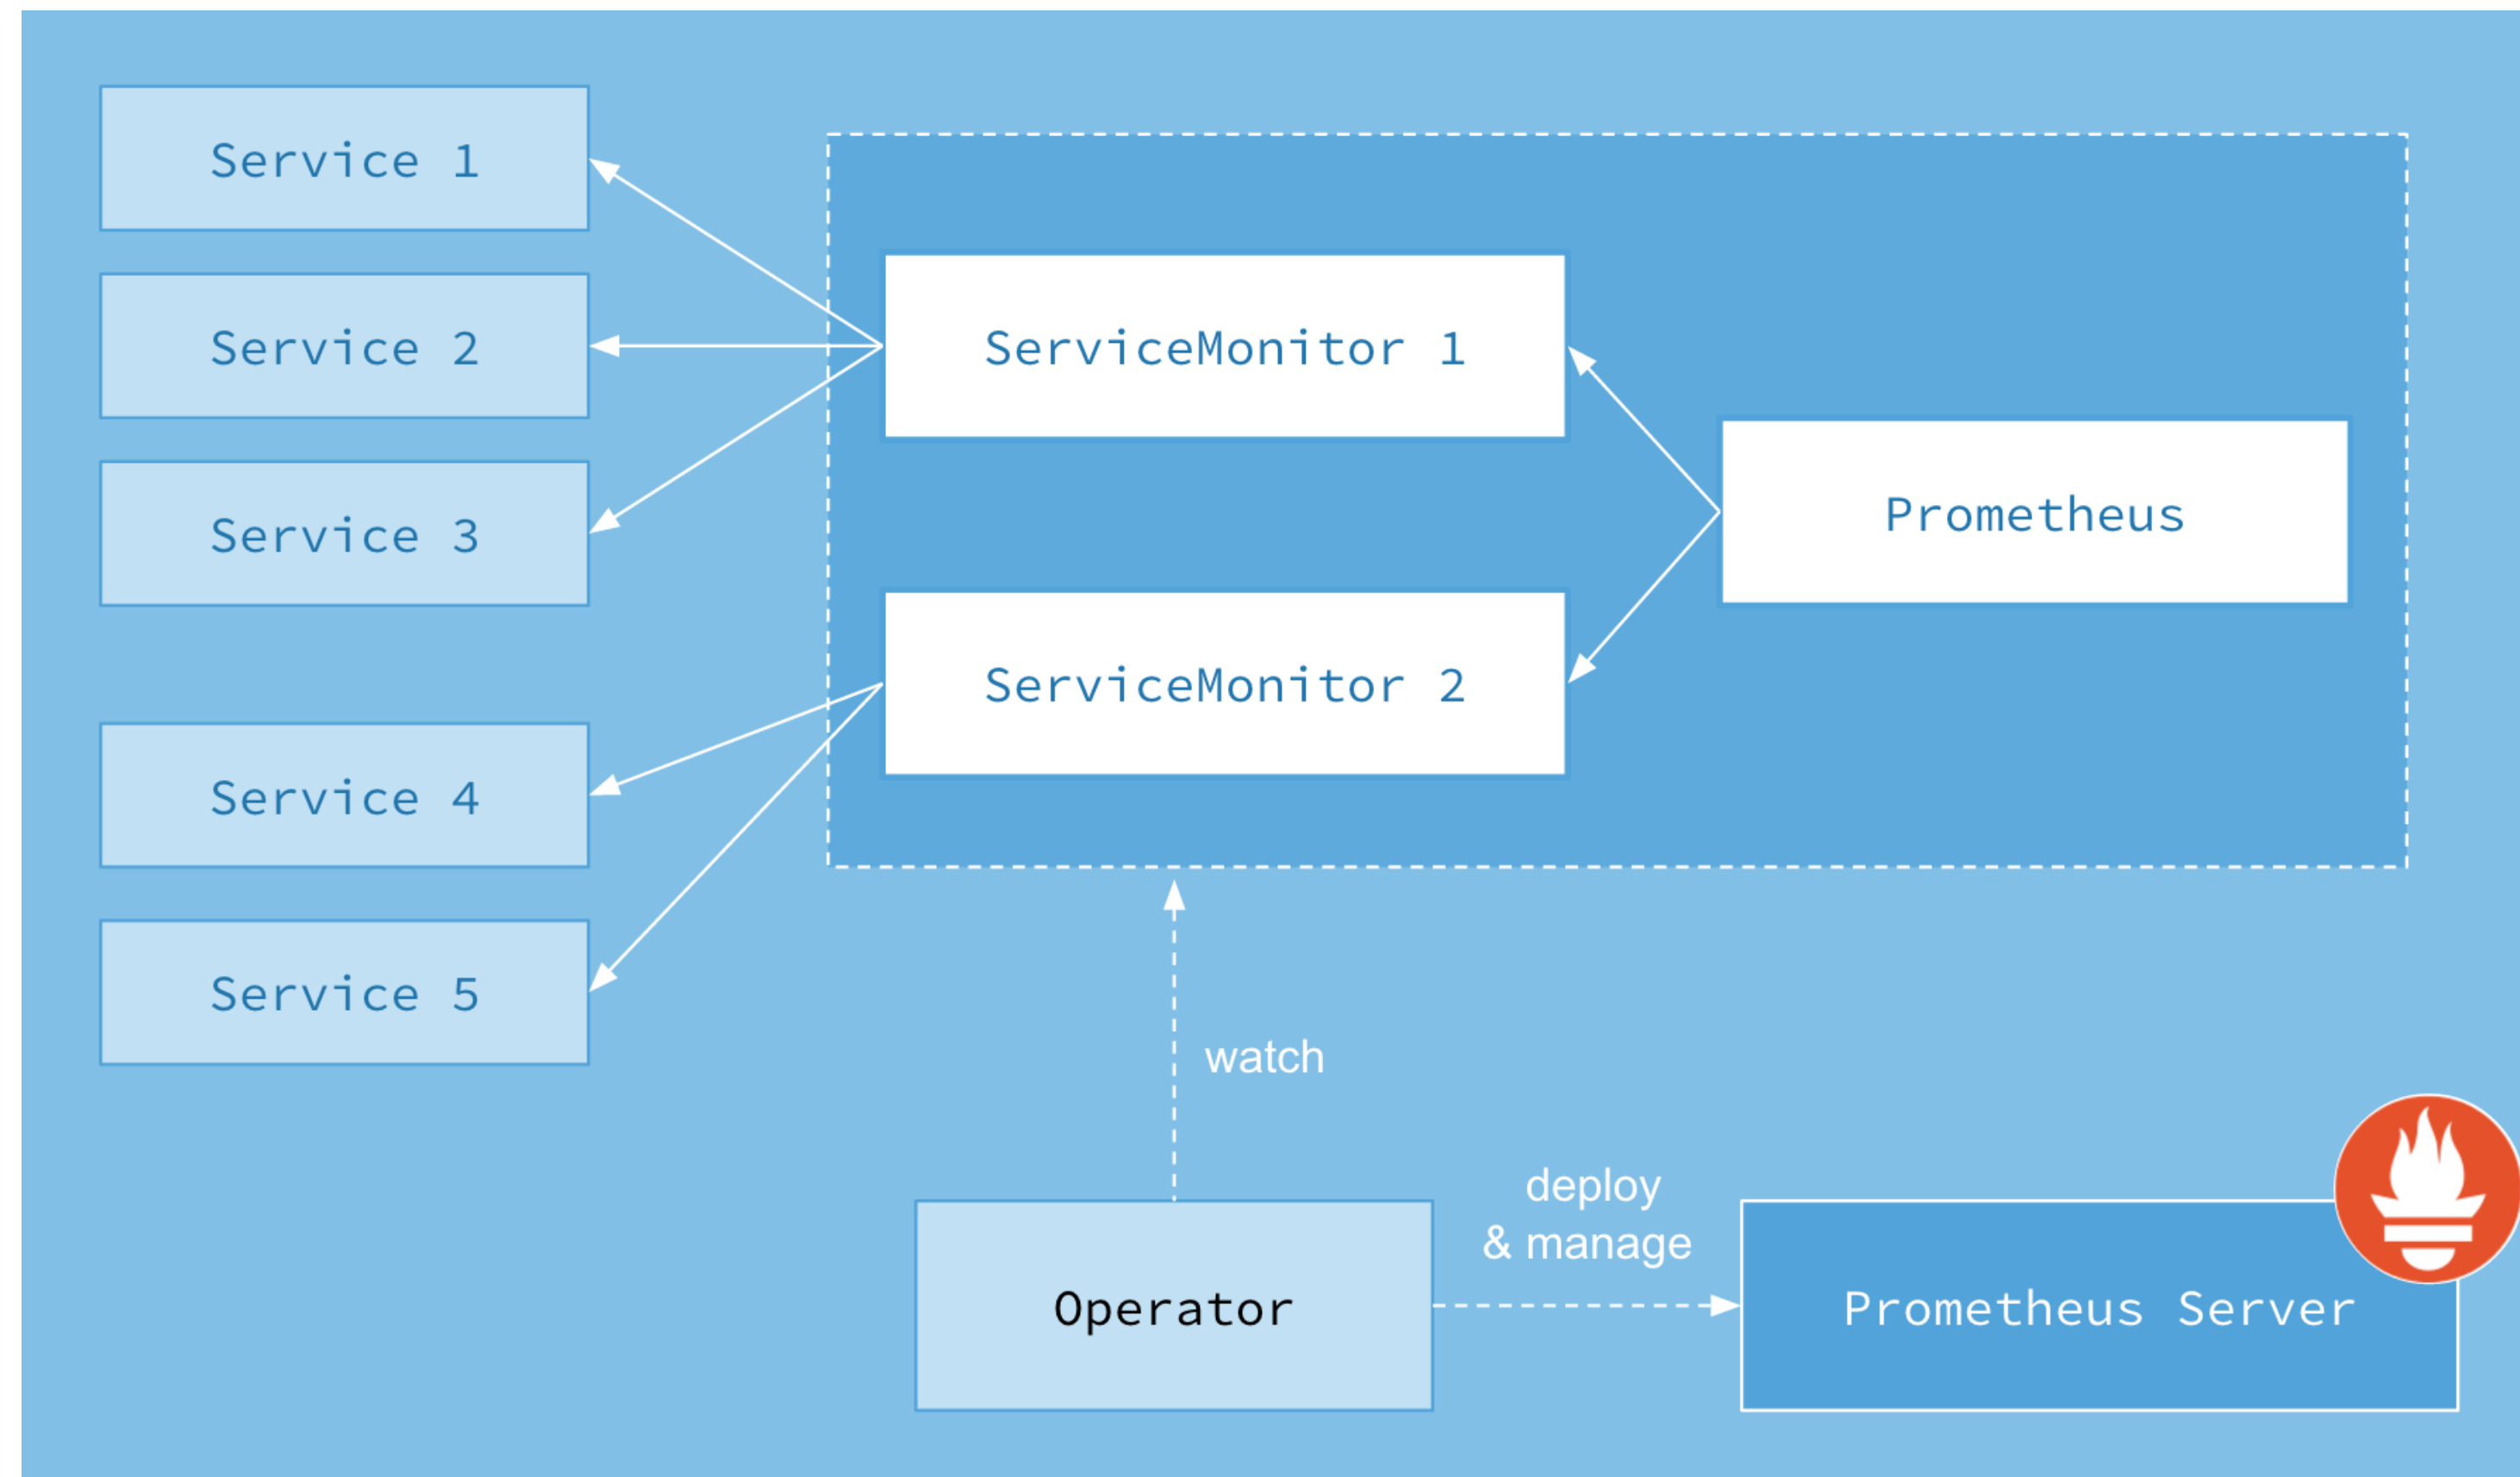
Kubernetes Operator



Kubernetes Operator



Prometheus Operator — Architecture Overview



ServiceMonitor

```
1  apiVersion: monitoring.coreos.com/v1
2  kind: ServiceMonitor
3  metadata:
4    labels:
5      app: prometheus-operator-alertmanager
6      chart: prometheus-operator-9.3.2
7      heritage: Helm
8      release: prometheus-operator
9      name: prometheus-operator-alertmanager
10     namespace: monitoring
11  spec:
12     endpoints:
13       - path: /metrics
14         port: web
15     namespaceSelector:
16       matchNames:
17         - monitoring
18     selector:
19       matchLabels:
20         app: prometheus-operator-alertmanager
21         release: prometheus-operator
```



ServiceMonitor

```
~ * wsp.production monitoring % k get servicemonitors.monitoring.coreos.com
```

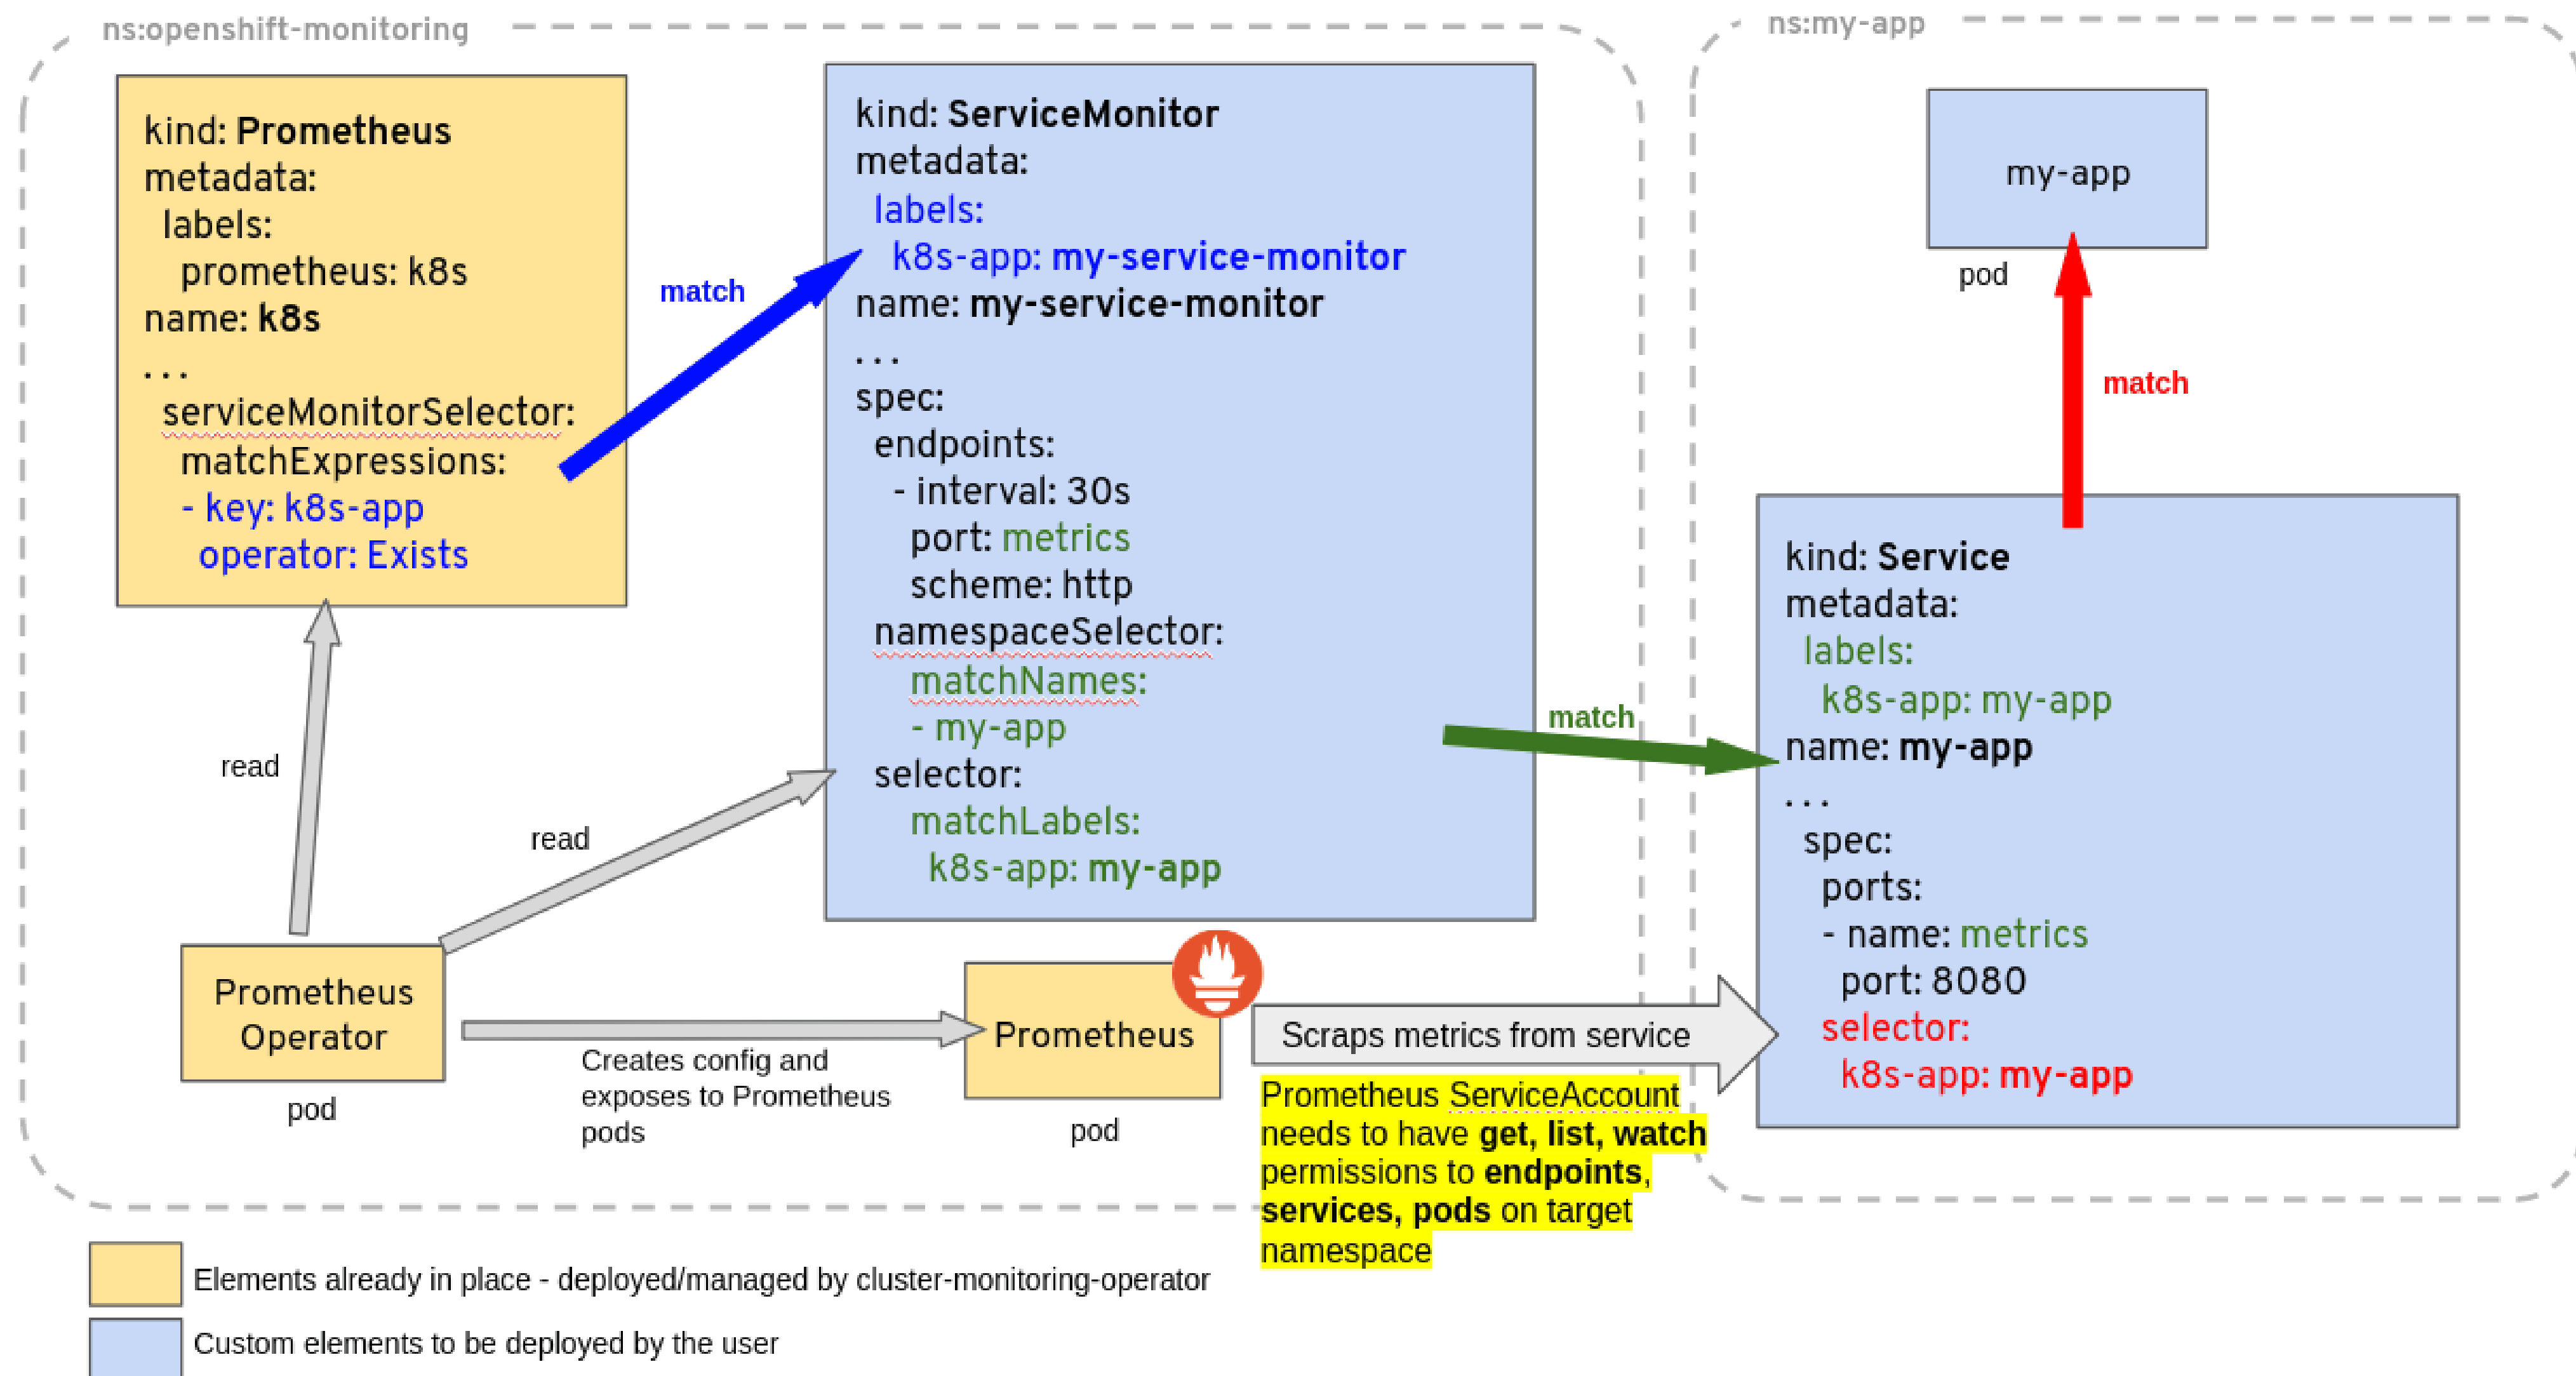
| NAME | AGE |
|---|------|
| prometheus-operator-alertmanager | 199d |
| prometheus-operator-apiserver | 199d |
| prometheus-operator-coredns | 199d |
| prometheus-operator-grafana | 199d |
| prometheus-operator-kube-controller-manager | 199d |
| prometheus-operator-kube-etcd | 199d |
| prometheus-operator-kube-proxy | 187d |
| prometheus-operator-kube-scheduler | 199d |
| prometheus-operator-kube-state-metrics | 199d |
| prometheus-operator-kubelet | 199d |
| prometheus-operator-node-exporter | 199d |
| prometheus-operator-operator | 199d |
| prometheus-operator-prometheus | 199d |



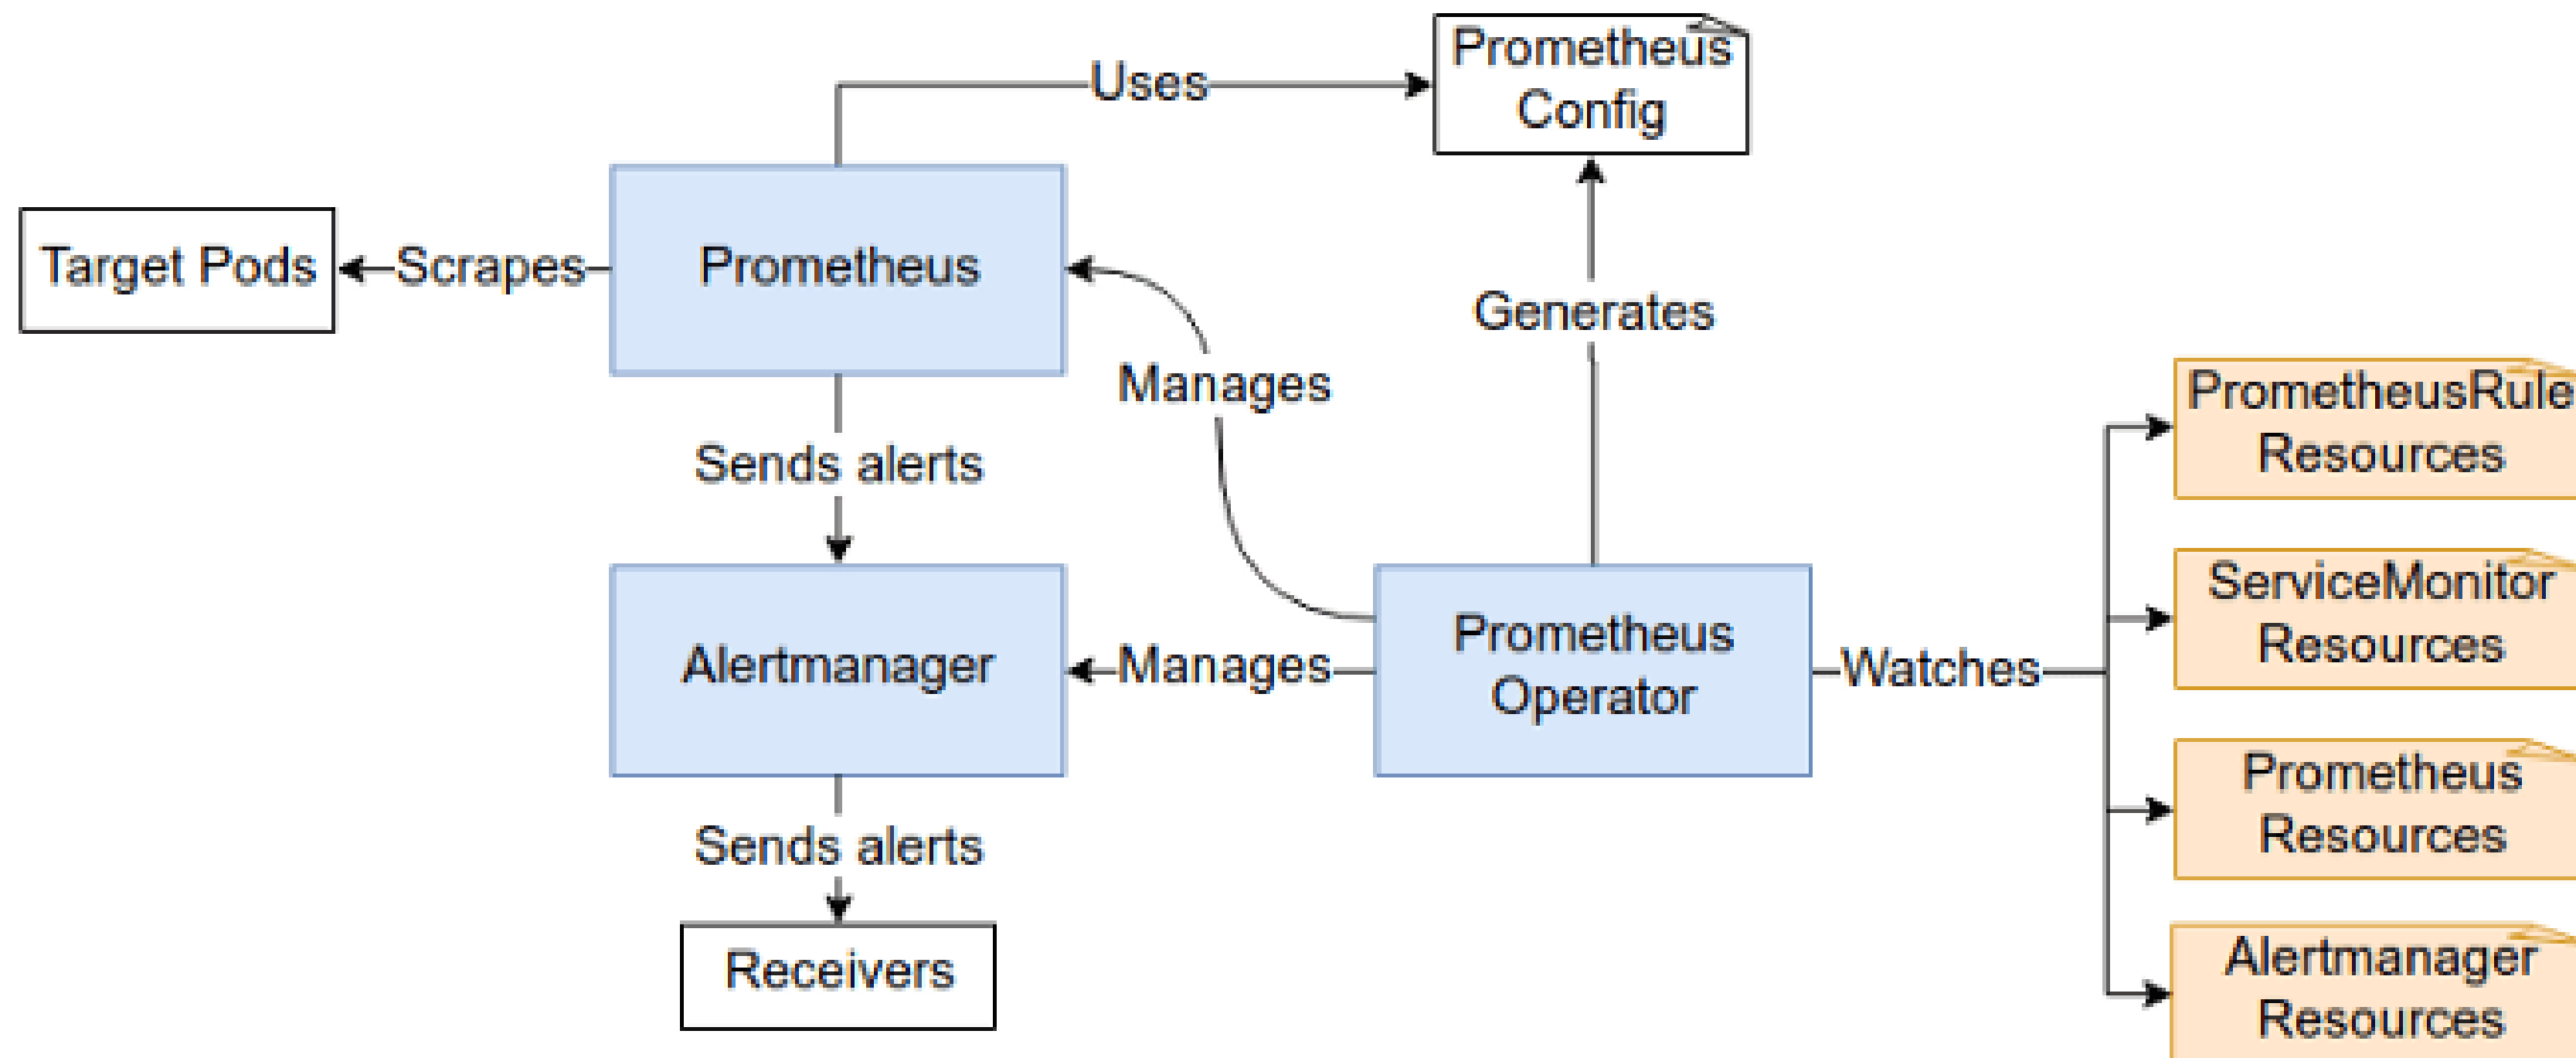
ServiceMonitor



ServiceMonitor



Prometheus Operator — how it works



PrometheusRules

```
1  apiVersion: monitoring.coreos.com/v1
2  kind: PrometheusRule
3  metadata:
4    annotations:
5      prometheus-operator-validated: "true"
6    labels:
7      app: prometheus-operator
8      chart: prometheus-operator-9.3.2
9      heritage: Helm
10     release: prometheus-operator
11     name: prometheus-operator-etcd
12     namespace: monitoring
13  spec:
14    groups:
15      - name: etcd
16        rules:
17          - alert: etcdInsufficientMembers
18            annotations:
19              message: 'etcd cluster "{{ $labels.job }}"': insufficient members ({{ $value
20                }}).'
21              expr: sum(up{job=~".*etcd.*"} == bool 1) by (job) < ((count(up{job=~".*etcd.*"})
22                by (job) + 1) / 2)
23              for: 3m
24            labels:
25              severity: critical
```


PrometheusRules — alerts out-of-the-box

```
~ * wsp.production monitoring % k get prometheusrules.monitoring.coreos.com
NAME                                     AGE
prometheus-operator-alertmanager       199d
prometheus-operator-blackbox           151d
prometheus-operator-cert-manager       151d
prometheus-operator-etcd               199d
prometheus-operator-general             199d
prometheus-operator-k8s                 199d
prometheus-operator-kube-apiserver-error 117d
prometheus-operator-kube-apiserver      199d
prometheus-operator-kube-prometheus-node-recording 199d
prometheus-operator-kube-scheduler      199d
prometheus-operator-kubernetes-absent   117d
prometheus-operator-kubernetes-apps     199d
prometheus-operator-kubernetes-resources 199d
prometheus-operator-kubernetes-storage  199d
prometheus-operator-kubernetes-system   199d
prometheus-operator-kubernetes-system-apiserver 199d
prometheus-operator-kubernetes-system-controller-manager 199d
prometheus-operator-kubernetes-system-kubelet 199d
prometheus-operator-kubernetes-system-scheduler 199d
```

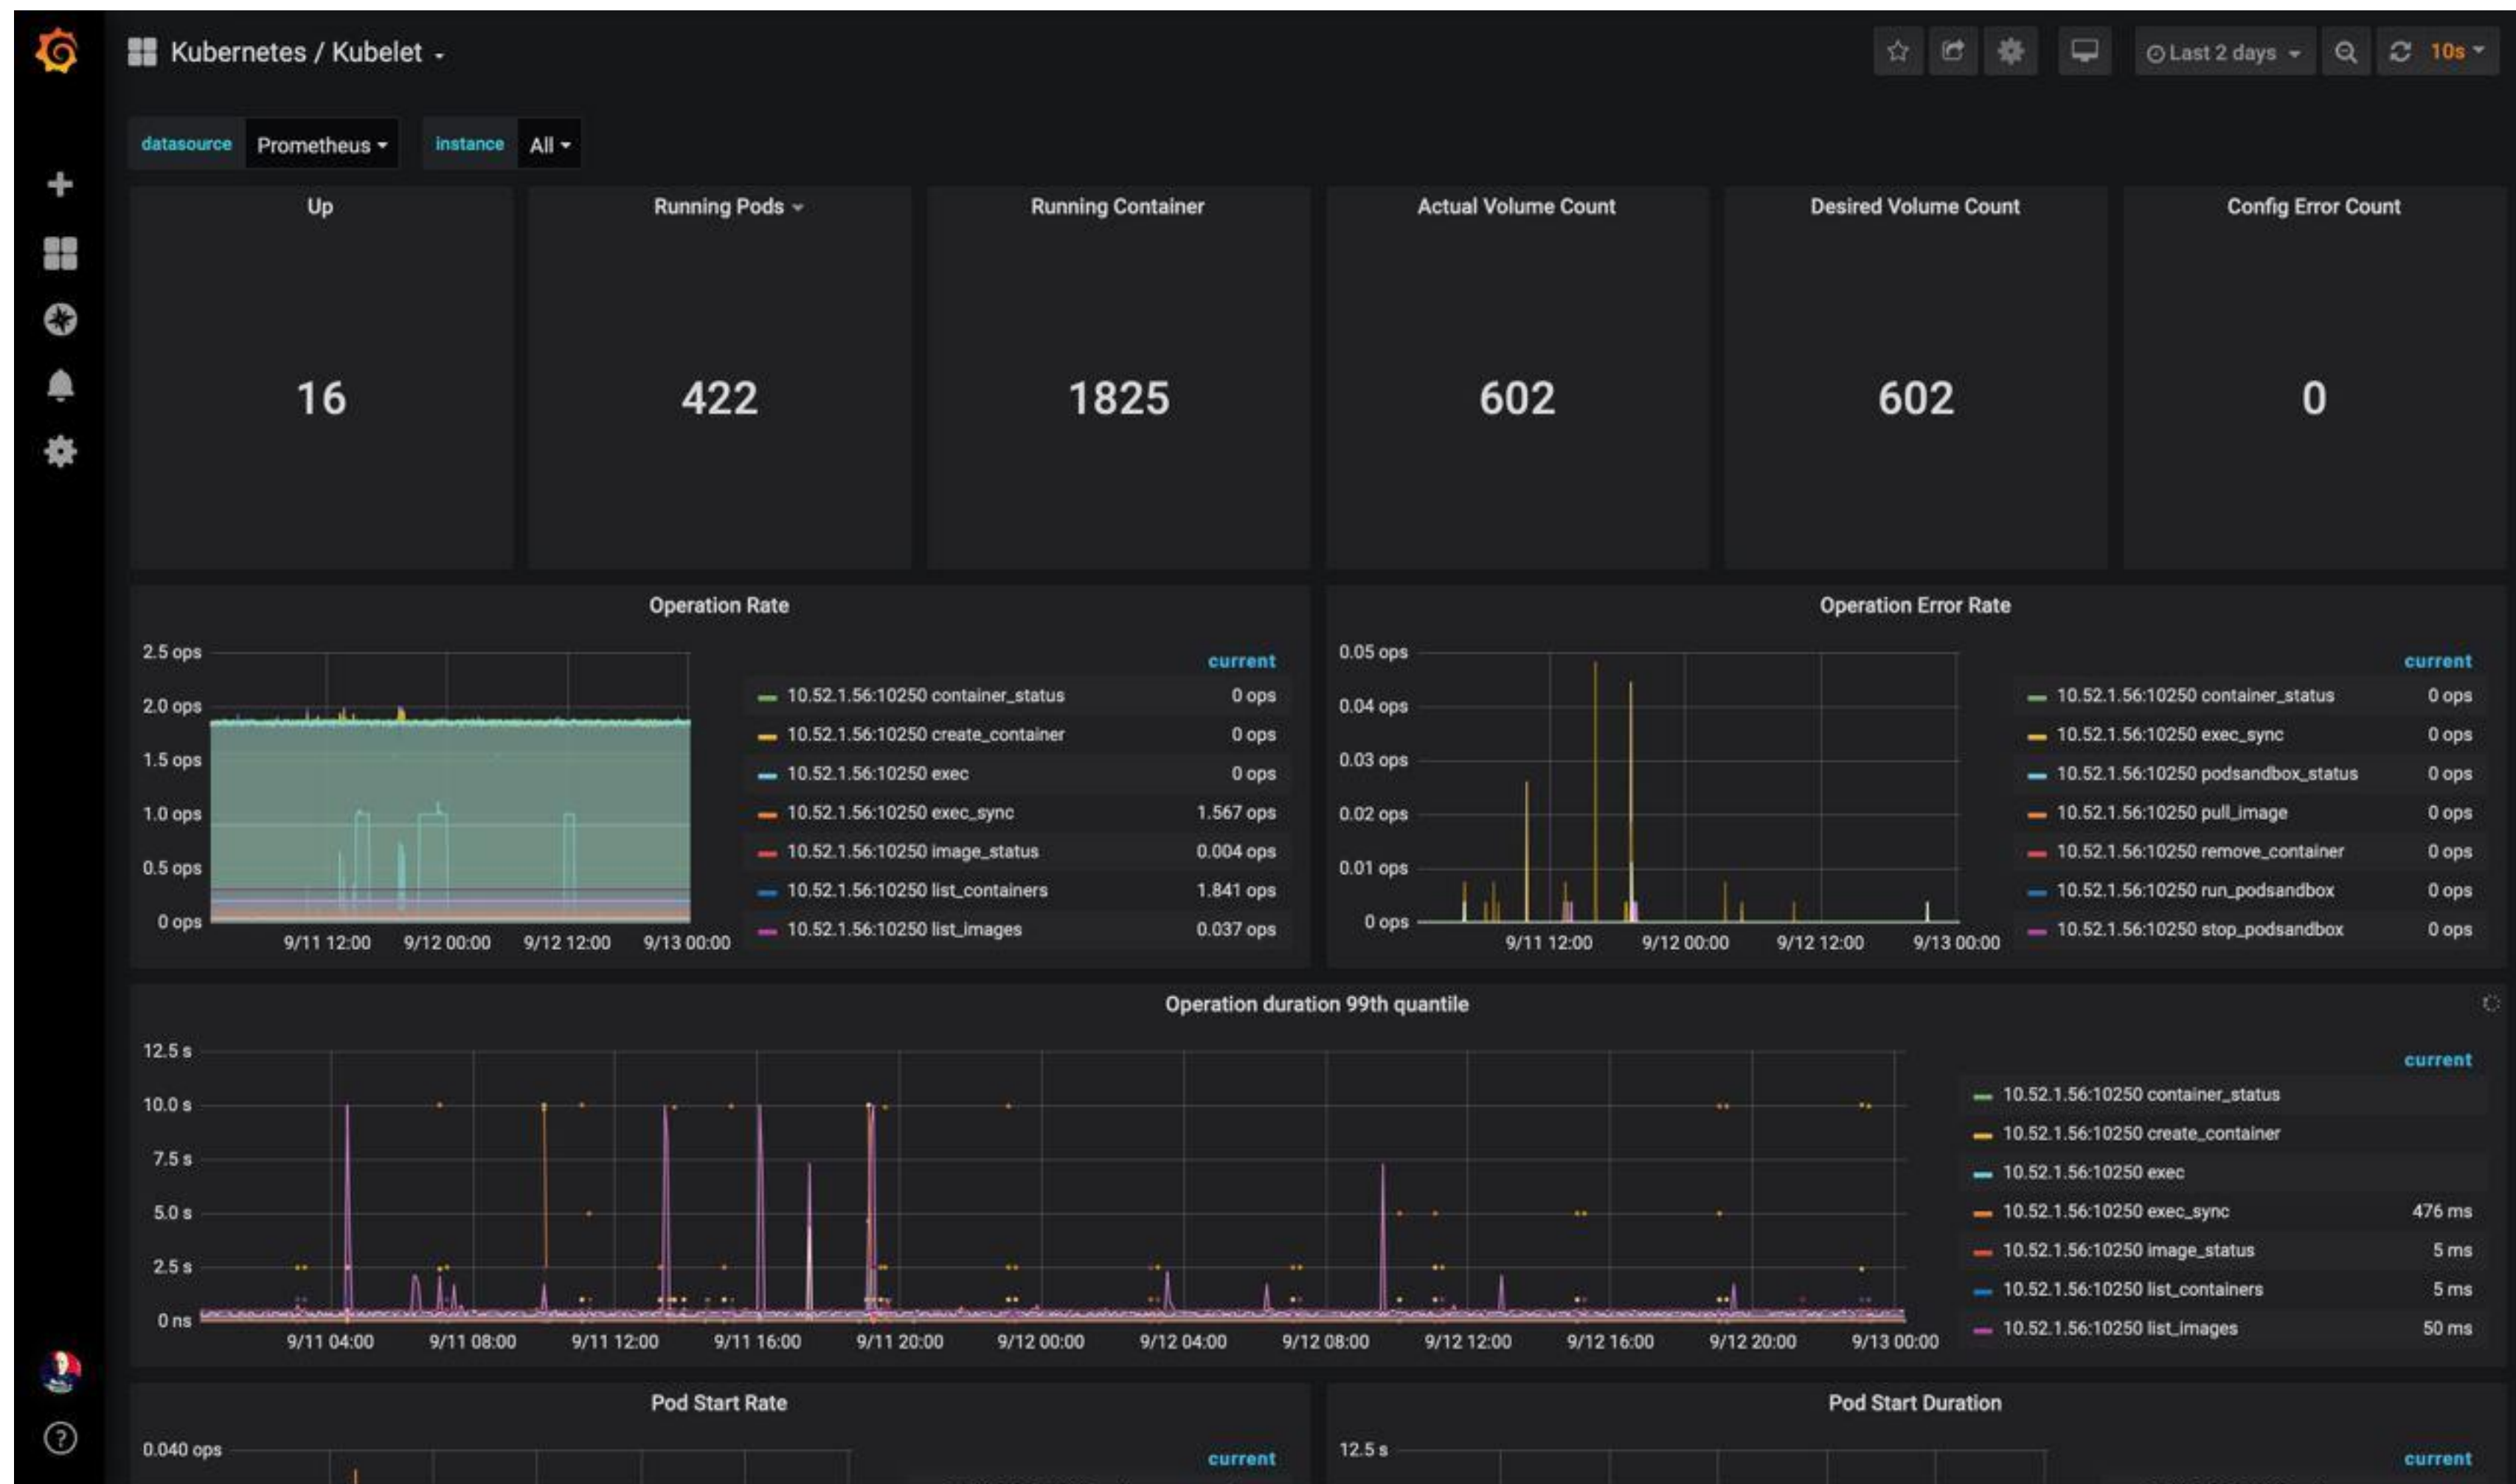

Kubernetes-mixin — all-in-one k8s monitoring

```
~ * wsp.production monitoring % k get configmaps | grep prometheus-operator
prometheus-operator-apiserver 1 201d
prometheus-operator-cluster-total 1 201d
prometheus-operator-controller-manager 1 201d
prometheus-operator-etcd 1 201d
prometheus-operator-grafana 2 201d
prometheus-operator-grafana-config-dashboards 1 201d
prometheus-operator-grafana-datasource 1 201d
prometheus-operator-grafana-test 1 201d
prometheus-operator-k8s-coredns 1 201d
prometheus-operator-k8s-resources-cluster 1 201d
prometheus-operator-k8s-resources-namespace 1 201d
prometheus-operator-k8s-resources-node 1 201d
prometheus-operator-k8s-resources-pod 1 201d
prometheus-operator-k8s-resources-workload 1 201d
prometheus-operator-k8s-resources-workloads-namespace 1 201d
prometheus-operator-kubelet 1 201d
prometheus-operator-namespace-by-pod 1 201d
prometheus-operator-namespace-by-workload 1 201d
prometheus-operator-node-cluster-rsrc-use 1 201d
prometheus-operator-node-rsrc-use 1 201d
prometheus-operator-nodes 1 201d
prometheus-operator-persistentvolumesusage 1 201d
prometheus-operator-pod-total 1 201d
prometheus-operator-pods 1 120d
prometheus-operator-prometheus 1 201d
prometheus-operator-proxy 1 189d
prometheus-operator-scheduler 1 201d
prometheus-operator-statefulset 1 201d
prometheus-operator-workload-total 1 201d
prometheus-prometheus-operator-prometheus-rulefiles-0 32 40d
```


Kubernetes-mixin — all-in-one k8s monitoring



Kubernetes-mixin — all-in-one k8s monitoring



Longterm storage Overview



Longterm storage Overview




Now we are here

- ~~• The way from Legacy monitoring systems to Modern~~
- ~~• Monitoring Problems~~
- ~~• Legacy Infrastructure Service Discovery~~
- ~~• Prometheus Operator as the Solution~~
- PaaS Alerting and Dashboards — Monitoring As A Code
- Upgrades and Incidents
- Conclusion

Deploy Helm Chart via Ansible





PROMETHEUS-OPERATOR VS KUBE-PROMETHEUS-STACK

DEPRECATED

Further development has moved to [prometheus-community/helm-charts](https://github.com/prometheus-community/helm-charts). The chart has been renamed [kube-prometheus-stack](#) to more clearly reflect that it installs the `kube-prometheus` project stack, within which Prometheus Operator is only one component.

PaaS Alerting and Dashboards — Monitoring As A Code



PaaS Alerting and Dashboards — Monitoring As A Code

```
- name: Applying prometheus crds
  k8s:
    definition:
      apiVersion: monitoring.coreos.com/v1
      kind: PrometheusRule
      metadata:
        labels:
          app: prometheus-operator
          generated: 'true'
          release: prometheus-operator
          name: "prometheus-operator-{{ (item.path | basename | splitext)[0] }}"
          namespace: monitoring
      spec:
        groups:
          - name: "{{ item.path | basename }}"
            rules: "{{ lookup('file', item.path) | from_yaml }}"
        loop: "{{ alerting_sources.files }}"
        loop_control:
          label: "{{ item.path }}"
```


PaaS Alerting and Dashboards — Monitoring As A Code

```
- alert: netdataMetricsUnreachable
  expr: netdata:reachability:bool == 0
  for: 5m
  labels:
    severity: critical
    routing_key: ops
  annotations:
    message: "{{ $labels.instance }} netdata metrics unreachable from prometheus for 5m."
    runbook: troubleshooting/node.md#netdataMetricsUnreachable
    dashboard_uid: 6bab49de
    panel_metric: netdata:reachability:bool{{"{"}}instance="{{ $labels.instance }}"{{"{"}}}}
```

PaaS Alerting and Dashboards — Monitoring As A Code

```
apiVersion: v1
kind: ConfigMap
metadata:
  labels:
    grafana_dashboard: "true"
    grafana_dashboard_custom: "true"
  namespace: monitoring
  name: grafana-dashboard-{{ (item.path | basename | splitext)[0] }}
data:
  {{ 'grafana-dashboard-%s' | format(item.path | basename) }}: >
  {{ lookup('file', item.path) | indent(6, first=True, blank=True) }}
```



Now we are here

- ~~• The way from Legacy monitoring systems to Modern~~
- ~~• Monitoring Problems~~
- ~~• Legacy Infrastructure Service Discovery~~
- ~~• Prometheus Operator as the Solution~~
- ~~• PaaS Alerting and Dashboards — Monitoring As A Code~~
- Upgrades and Incidents
- Conclusion

Upgrades and Incidents



Upgrades and Incidents

- ETCD endpoints — не работают из коробки

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade
- Удаление ServiceMonitor при удалении helm-релиза

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade
- Удаление ServiceMonitor при удалении helm-релиза
- Неоптимальные параметры ServiceMonitor

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade
- Удаление ServiceMonitor при удалении helm-релиза
- Неоптимальные параметры ServiceMonitor
- Обновление конфигурации Prometheus — instance restart

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade
- Удаление ServiceMonitor при удалении helm-релиза
- Неоптимальные параметры ServiceMonitor
- Обновление конфигурации Prometheus — instance restart
- Отсутствие метрик (nodata) — тоже метрика

Upgrades and Incidents

- ETCD endpoints — не работают из коробки
- Изменение CRD при апгрейде = fail upgrade
- Удаление ServiceMonitor при удалении helm-релиза
- Неоптимальные параметры ServiceMonitor
- Обновление конфигурации Prometheus — instance restart
- Отсутствие метрик (nodata) — тоже метрика
- Kube-apiserver умирает при 4GB при деплое helm-chart

Now we are here


- ~~• The way from Legacy monitoring systems to Modern~~
- ~~• Monitoring Problems~~
- ~~• Legacy Infrastructure Service Discovery~~
- ~~• Prometheus Operator as the Solution~~
- ~~• PaaS Alerting and Dashboards — Monitoring As A Code~~
- ~~• Upgrades and Incidents~~
- Conclusion

Выводы

- “Админы” больше не bottleneck
- Разработчики самостоятельно доставляют бизнес-метрики
- Решение легко встраивается в любой k8s pipeline
- Легко ложится на infrastructure as a code



Q/A and Contacts



sshcherbakov@plesk.com

twitter.com/stasian

t.me/xSTASiANx